

RESEARCH ARTICLE

Open Access



# KekuleScope: prediction of cancer cell line sensitivity and compound potency using convolutional neural networks trained on compound images

Isidro Cortés-Ciriano\* and Andreas Bender

## Abstract

The application of convolutional neural networks (ConvNets) to harness high-content screening images or 2D compound representations is gaining increasing attention in drug discovery. However, existing applications often require large data sets for training, or sophisticated pretraining schemes. Here, we show using 33 IC<sub>50</sub> data sets from ChEMBL 23 that the in vitro activity of compounds on cancer cell lines and protein targets can be accurately predicted on a continuous scale from their Kekulé structure representations alone by extending existing architectures (AlexNet, DenseNet-201, ResNet152 and VGG-19), which were pretrained on unrelated image data sets. We show that the predictive power of the generated models, which just require standard 2D compound representations as input, is comparable to that of Random Forest (RF) models and fully-connected Deep Neural Networks trained on circular (Morgan) fingerprints. Notably, including additional fully-connected layers further increases the predictive power of the ConvNets by up to 10%. Analysis of the predictions generated by RF models and ConvNets shows that by simply averaging the output of the RF models and ConvNets we obtain significantly lower errors in prediction for multiple data sets, although the effect size is small, than those obtained with either model alone, indicating that the features extracted by the convolutional layers of the ConvNets provide complementary predictive signal to Morgan fingerprints. Lastly, we show that multi-task ConvNets trained on compound images permit to model COX isoform selectivity on a continuous scale with errors in prediction comparable to the uncertainty of the data. Overall, in this work we present a set of ConvNet architectures for the prediction of compound activity from their Kekulé structure representations with state-of-the-art performance, that require no generation of compound descriptors or use of sophisticated image processing techniques. The code needed to reproduce the results presented in this study and all the data sets are provided at <https://github.com/isidro/kekulescope>.

**Keywords:** Convolutional Neural Networks, Deep learning, Drug discovery, Bioactivity modelling, Random Forest, Cytotoxicity, ChEMBL

## Introduction

Cultured cancer cell lines are limited disease models in that they do not recapitulate the tumor microenvironment nor interactions with the immune system [1–6], fundamental properties of cellular organization are

altered in culture [7], and their response to anticancer drugs is affected by both assay heterogeneity [8] and genomic alterations acquired in vitro [9]. However, cancer cell lines still represent versatile models to study fundamental aspects of cancer biology [10, 11], and the genomic determinants of drug response [3, 12–14]. Hence, the development of computational methods to harness the large amount of in vitro cell line sensitivity data collected to date to unravel the underlying

\*Correspondence: [isidrolauscher@gmail.com](mailto:isidrolauscher@gmail.com)

Centre for Molecular Informatics, Department of Chemistry, University of Cambridge, Lensfield Road, Cambridge CB2 1EW, UK



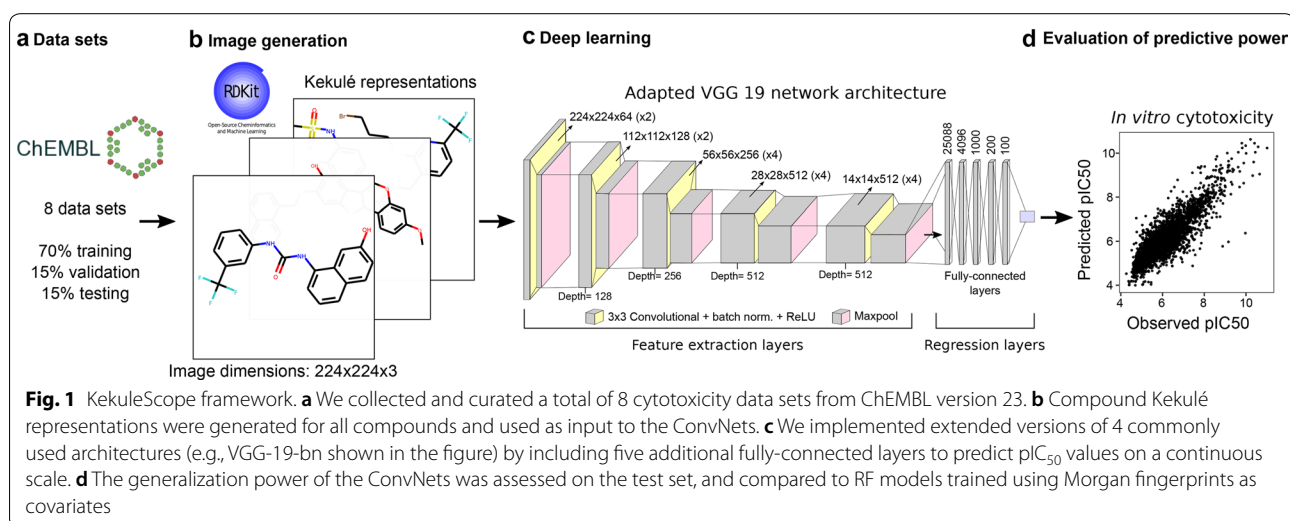
molecular mechanisms mediating drug activity and identify novel biomarkers for drug response is an area of intense research [14–20].

Whereas existing computational tools to model in vitro compound activity mostly rely on established algorithms (e.g., Random Forest or Support Vector Machines), the utilization of deep learning in drug discovery is gaining momentum, a trend that is only expected to increase in the coming years [21]. Deep learning techniques have been already applied in numerous drug discovery tasks, including toxicity modelling [22, 23], bioactivity prediction [24–30], and de novo drug design [31–34], among others. Most of these studies have utilized feedforward neural networks consisting of multiple fully-connected layers trained on one of the many compound descriptors developed over the last > 30 years in the cheminformatics field [27, 35]. However, the high performance of convolutional neural networks (ConvNets) [36–38], a type of neural network developed for image recognition tasks, in finding complex high-dimensional relationships in diverse image data sets is fostering their application in drug discovery [21, 39, 40].

ConvNets consist of two sets of layers (Fig. 1): (1) the convolutional layers, which extract features from the input images, and (2) the classification/regression layers, which are generally fully-connected layers that output one value for each of the tasks being modelled. A major advantage of ConvNets is that the extraction of features is performed on a fully automatic and data-driven fashion, thus not requiring to engineer feature selection or image preprocessing filters beforehand [33, 39, 41, 42]. Today, convolutional neural networks are applied to diverse image recognition tasks in healthcare and biomedicine

[43–46]. An obvious critical element for the application of ConvNets is the availability of images for training, or the ability to formulate the modelling task of interest as an image classification problem. An illustrative example of the latter is DeepVariant [47], a recently published algorithm that uses images of sequencing read pileups as input to detect small indels and single-nucleotide variants, instead of assigning probabilities to each of the genotypes supported by the data using statistical modelling, as has been the standard approach for years.

In drug discovery, applications of ConvNets include elucidation of the mechanism of action of small molecules and their bioactivity profiles from high-content screening images [48–50], and modelling in vitro assay endpoints using 2D representations of compound structures, termed “compound images”, as input [23, 41, 51–53]. Efforts to model compound activity using ConvNets trained on compound images were spearheaded by Goh et al., who developed *Chemception* [51, 54], a ConvNet based on the Inception-ResNet v2 architecture [55]. The performance of *Chemception* was compared to multi-layer perceptron deep neural networks trained on circular fingerprints in three tasks: prediction of free energy of solvation (632 compounds; regression), inhibition of HIV replication (41,193; binary classification), and compound toxicity using data from the “Toxicology in the 21st Century” (Tox21) project (8014; multi-task binary classification) [56]. *Chemception* slightly outperformed the multi-layer perceptron networks except for the TOX21 task. In a follow-up study, the same group introduced *ChemNet* [51], a learning strategy that consists of pre-training a ConvNet (e.g., *Chemception* [54]) using a large set of compounds (1.7 M) to predict their physico-chemical properties (e.g., logP, which represents an easy



task) in order to learn general features related to chemistry from the images. Subsequently, the trained networks were applied to model smaller data sets using transfer learning. Although such an approach led to higher performance than *Chemception*, a major disadvantage thereof is that it requires the initial training of the network on a large set of compounds, which is computationally demanding. More recently, Fernández et al. proposed *Toxic Colors*, a framework to classify toxic compounds from the TOX21 data set using compound images as input [23]. Although these studies have paved the way for the application of ConvNets to model the bioactivity of compounds using their images as input, a comprehensive analysis of ConvNet architectures with a reduced computational footprint to model cancer cell line sensitivity on a continuous scale and comparison against the state of the art is still missing. Moreover, whether the combination of models trained on widely-used compound descriptors (e.g., circular fingerprints) and ConvNets trained using compound images leads to increased predictive power remains to be studied.

Here, we introduce *KekuleScope*, a flexible framework for modelling the bioactivity of compounds on a continuous scale from their Kekulé structure representation using ConvNets pretrained to model unrelated image classification tasks. We demonstrate using 8 cytotoxicity data sets and in vitro IC<sub>50</sub> data for 25 diverse protein targets extracted from ChEMBL version 23 (Table 1) that compound images convey enough predictive power to build robust models using ConvNets. Instead of using networks pretrained on compound images [51], we show that widely-used architectures developed for unrelated image classification tasks (AlexNet [57], DenseNet-201 [58], ResNet152 [59] and VGG-19 [60]) are versatile enough to generate robust predictions across a dynamic range of bioactivity values using compound images as input. Moreover, comparison with Random Forest models and Deep Neural Networks (DNN) trained on circular

fingerprints (Morgan fingerprints [61, 62]) reveals that ConvNets trained using compound images lead to comparable predictive power on the test set. In addition, combining RF and ConvNet predictions into model ensembles often leads to increased model performance, suggesting that the features extracted by the convolutional layers of the networks provide complementary information to Morgan fingerprints. Therefore, our work presents a novel framework for the prediction of compound activity that requires minimal deep learning architecture design, processing of chemical structures and no descriptor choice, and that leads to improved predictive power over the state of the art in our validation on 8 cancer cell line sensitivity and 25 in vitro potency data sets.

## Methods

### Data collection and curation

We gathered cytotoxicity IC<sub>50</sub> data for 8 cancer cell lines and 25 protein targets from ChEMBL database version 23 using the *chembl\_webresource\_client* Python module [63–65]. To gather high-quality bioactivity data sets, we only kept IC<sub>50</sub> values for small molecules that satisfied the following stringent filtering criteria [8]: (1) activity unit equal to “nM”, and (2) activity relationship equal to ‘=’. The average pIC<sub>50</sub> value was calculated when multiple IC<sub>50</sub> values were annotated for the same compound-cell line or compound-protein pair. IC<sub>50</sub> values were modeled in a logarithmic scale ( $pIC_{50} = -\log_{10} IC_{50} [M]$ ). We selected the data sets on a purely data-driven fashion, as these are the protein targets with the highest number of IC<sub>50</sub> values available (after applying the stringent filtering and data curation criteria specified above). As for the cell lines, we selected these 8 on the basis of data availability as well, and because they are commonly used in preclinical drug discovery. Further information about the data sets is given in Tables 1 and 2. All data sets used in this study are available at <https://github.com/isidro/kekulescope>.

**Table 1** Cell line data sets used in this study

Cell line	Description	ChEMBL cell ID	Cellosaurus ID	Organism of origin	Number of bioactivity data points
A2780	Ovarian carcinoma cells	CHEMBL3308421	CVCL_0134	<i>Homo sapiens</i>	2255
CCRF-CEM	T-cell leukemia	CHEMBL3307641	CVCL_0207	<i>Homo sapiens</i>	3047
DU-145	Prostate carcinoma	CHEMBL3308034	CVCL_0105	<i>Homo sapiens</i>	2512
HCT-15	Colon adenocarcinoma cells	CHEMBL3307945	CVCL_0292	<i>Homo sapiens</i>	994
KB	Squamous cell carcinoma	CHEMBL3307959	CVCL_0372	<i>Homo sapiens</i>	2731
LoVo	Colon adenocarcinoma cells	CHEMBL3307691	CVCL_0399	<i>Homo sapiens</i>	1120
PC-3	Prostate carcinoma cells	CHEMBL3307570	CVCL_0035	<i>Homo sapiens</i>	4294
SK-OV-3	Ovarian carcinoma cells	CHEMBL3307746	CVCL_0532	<i>Homo sapiens</i>	1589

### Molecular representation

We standardized all chemical structures to a common representation scheme using the Python module *standardizer* (<https://github.com/flatkinson/standardizer>). Entries containing inorganic elements were entirely removed from the data sets, and the largest fragment was kept to remove counterions and solvents. We note that, although imperfect, removing counterions is a standard procedure in the field [66, 67]. In addition, salts are not generally well-handled by descriptor calculation software, and hence, filtering them out is generally preferred [68].

Kekulé structure representations for all compounds (i.e., ‘compound images’) in Scalable Vector Graphics (SVG) format were generated from the compound structures in SDF format using the RDKit function *MolsToGridImage* and default parameter values. SVG images were then converted to Portable Network Graphics (PNG) format using the programme *convert* (version ImageMagick 6.7.8-9 2016-03-31 Q16; <http://www.imagemagick.org>) and resized to 224 × 224 pixels using a density (−d argument) of 800. The code needed to reproduce the

results presented in this study is provided at <https://github.com/isidro/kekulescope>. To represent molecules for subsequent model generation based on fingerprints, we computed circular Morgan fingerprints [61] for all compounds using RDKit (release version 2013.03.02) [69]. The radius was set to 2 and the fingerprint lengths to 128, 256, 512, 1024 and 2048.

### Machine learning

#### Data splitting

The data sets were randomly split into a training (70% of the data), validation (15%), and test set (15%). For each data set, the training set was used to train the ConvNets, the validation set served to monitor their predictive power during the training phase, and the test set served to assess their predictive power on unseen data after the ConvNets were trained.

#### Convolutional neural network architectures and training

ConvNets pretrained on the ImageNet [70] data set were downloaded using the Python library Pytorch [71]. The structure of the classification layer(s) in each of the

**Table 2 Protein target data sets used in this study**

Target preferred name	Target abbreviation	Uniprot ID	ChEMBL ID	Number of bioactivity data points
Alpha-2a adrenergic receptor	A2a	P08913	CHEMBL1867	203
Tyrosine-protein kinase ABL	ABL1	P00519	CHEMBL1862	773
Acetylcholinesterase	Acetylcholinesterase	P22303	CHEMBL220	3159
Androgen Receptor	Androgen	P10275	CHEMBL1871	1290
Serine/threonine-protein kinase Aurora-A	Aurora-A	O14965	CHEMBL4722	2125
Serine/threonine-protein kinase B-raf	B-raf	P15056	CHEMBL5145	1730
Cannabinoid CB1 receptor	Cannabinoid	P21554	CHEMBL218	1116
Carbonic anhydrase II	Carbonic	P00918	CHEMBL205	603
Caspase-3	Caspase	P42574	CHEMBL2334	1606
Thrombin	Coagulation	P00734	CHEMBL204	1700
Cyclooxygenase-1	COX-1	P23219	CHEMBL221	1343
Cyclooxygenase-2	COX-2	P35354	CHEMBL230	2855
Dihydrofolate reductase	Dihydrofolate	P00374	CHEMBL202	584
Dopamine D2 receptor	Dopamine	P14416	CHEMBL217	479
Norepinephrine transporter	Ephrin	P23975	CHEMBL222	1740
Epidermal growth factor receptor erbB1	erbB1	P00533	CHEMBL203	4868
Estrogen receptor alpha	Estrogen	P03372	CHEMBL206	1705
Glucocorticoid receptor	Glucocorticoid	P04150	CHEMBL2034	1447
Glycogen synthase kinase-3 beta	Glycogen	P49841	CHEMBL262	1757
HERG	HERG	Q12809	CHEMBL240	5207
Tyrosine-protein kinase JAK2	JAK2	O60674	CHEMBL2971	2655
Tyrosine-protein kinase LCK	LCK	P06239	CHEMBL258	1352
Monoamine oxidase A	Monoamine	P21397	CHEMBL1951	1379
Mu opioid receptor	Opioid	P35372	CHEMBL233	840
Vanilloid receptor	Vanilloid	Q8NER1	CHEMBL4794	1923

architectures used was modified to output a single value, corresponding to compound  $pIC_{50}$  values in this case, by removing the softmax transformation of the last fully connected layer (which is used in classification tasks to output class scores in the 0–1 range). The Root Mean Squared Error (RMSE) value on the validation set was used as the loss function during the training phase of the ConvNets, and to compare the predictive power of RF, fully-connected neural networks, and ConvNets on the test set. We performed grid search to find the optimal combination of parameter values for all networks. The parameter values considered are listed in Table 3.

We generated an extended version of each architecture by including five fully-connected layers, consisting of 4096, 1000, 200 and 100 neurons (Fig. 1). Thus, for each architecture we implemented two regression versions, one containing one fully-connected layer, and a second one containing five fully-connected layers (abbreviated from now on as “extended”). The feature extraction layers were not modified.

In cases where the data sets were augmented, the following transformations were applied (as implemented in the Pytorch [71] library): (1) 180° rotation about the vertical axis (function *transforms.RandomHorizontalFlip*); (2) 180° rotation about the horizontal axis (*transforms.RandomVerticalFlip*); and (3) random 90° rotation (*transforms.RandomRotation*). In the three cases, each transformation was applied at every epoch during the training phase with a 50% chance. Thus, in some cases a set of the images might remain intact depending on this sampling step during a given epoch.

We used stochastic Gradient Descent algorithm with Nesterov momentum [72] to train all networks, which was set to 0.9 and kept constant during the training phase [72]. The parameters for all layers, including the convolutional and regression layers, were optimized during the training phase. Networks were allowed to evolve over 600 epochs. The networks were allowed to evolve over 600 epochs because we did not observe an increase in

predictive power in our initial experiments if we trained for more epochs. Given the high computational cost associated to training these models we decided that 600 epochs represent an appropriate trade-off between computational cost and predictive power (see Fig. 2).

To reduce the chance of overfitting, we used (1) early stopping, i.e., the training phase was stopped if the validation loss did not decrease after 250 epochs, and (2) 50% dropout [27, 73] in the five fully-connected layers (labelled as “Regression layers” in Fig. 1) in the extended versions of the architectures considered. The training phase was divided into cycles of 200 epochs, throughout which the learning rate was annealed and set back to its original value at the beginning of the next cycle. The learning rate was decreased by 90% or 40% every 10 or 25 epochs (decay rates of 0.1 and 0.6, respectively; Table 3).

#### Fully-connected deep neural networks (DNN)

DNN were trained using the Python library Pytorch [71] as previously described [74]. Briefly, we defined three hidden layers, composed of 60, 20, and 10 nodes, respectively, and used 10% dropout in the three hidden layers [27, 73]. The RMSE value on the validation set was used as the loss function during training. The training data were processed in batches of size equal to 15% of the number of instances. Rectified linear unit (ReLU) activation [27], and stochastic gradient descent with Nesterov momentum, which was set to 0.9 and kept constant during the training phase [72], were used to train all networks. The networks were allowed to evolve over 2000 epochs, and early stopping was performed in cases where the validation loss did not decrease after 200 consecutive epochs. We used 2000 epochs because this number was long enough to reach convergence of the networks. We note that the computational cost associated to training fully-connected networks using Morgan fingerprints is much smaller than the computational footprint of image-based models, which permitted us to train longer. We note that longer training times for networks using Morgan fingerprints can only result in an advantage for these over the image-based ones. The fact that the performance of fully-connected networks trained on Morgan fingerprints and networks trained on images is comparable indicates that we are not biasing our results in favor of the image-based models.

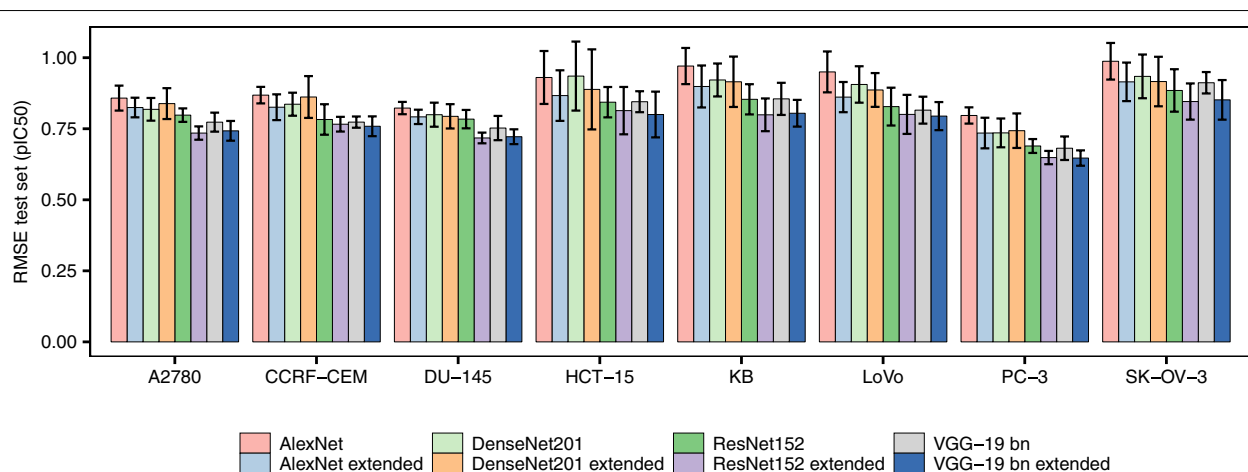
#### Random Forest (RF)

RF models based on Morgan fingerprint representations, which were calculated as described above, were generated using the Python library scikit-learn [75]. Default parameter values were used except for the number of trees, which was set to 100 because higher values do not

**Table 3** Parameters tuned during the training phase using grid search

Parameter	Values evaluated
Learning rate (Lr)	{0.1, 0.01, 0.001, 0.005, 0.001, 0.0001}
Decay rate	{0.1, 0.6}
Annealing rate step	{10, 25}
Data augmentation (augmentation)	{Yes: 1, No: 0}
Batch size (batch)	{4, 16, 32}

The names in parentheses indicate the parameter name abbreviation used in the main text and figures



**Fig. 2** Benchmarking the predictive power of ConvNet architectures on cytotoxicity data sets. Mean RMSE values ( $\pm$  standard deviation) on the test set across ten runs for each of the ConvNet architectures explored in this study (AlexNet [57], DenseNet-201 [58], ResNet152 [59] and VGG-19 [60]). Overall, all architectures enabled the generation of models with high predictive power on the test set, with RMSE values in the 0.65–0.96 pIC<sub>50</sub> range. However, the extended versions of these architectures that we designed by including 5 fully-connected layers (see Fig. 1) constantly led to increased predictive power on the test set

generally increase model performance when modelling bioactivity data sets [15, 76]. Identical data splits were used to train the ConvNets, DNN and the RF models.

### Experimental design

To compare the predictive power of the ConvNets to model cell line sensitivity in a robust statistical manner we designed a balanced fixed-effect full-factorial experiment with replications [77]. The following factors were considered:

1. *Data set*: 8 cytotoxicity data sets (Table 1).
2. *Model*: 8 convolutional network architectures.
3. *Batch size (Batch)*: number of compound images processed in each batch during the training phase.
4. *Data Augmentation (Augmentation)*: binary variable indicating whether data augmentation was applied during the training phase.

We implemented the following linear model to study this factorial design:

$$\begin{aligned}
 pIC_{50} = & \text{Data set}_i + \text{Model}_j + \text{Batch}_k \\
 & + \text{Augmentation}_l + \mu_0 + \varepsilon_{i,j,k,l,m} \\
 (i \in & \{1, \dots, N_{\text{datasets}} = 8\}; \\
 j \in & \{1, \dots, N_{\text{models}} = 8\}; \\
 k \in & \{1, \dots, N_{\text{batch sizes}} = 3\}; \\
 l \in & \{1, \dots, N_{\text{augmentation}} = 2\}; \\
 m \in & \{1, \dots, N_{\text{repetitions}} = 10\})
 \end{aligned} \quad (1)$$

where the factors *Data set*<sub>*i*</sub>, *Model*<sub>*j*</sub>, *Batch*<sub>*k*</sub>, *Augmentation*<sub>*l*</sub>, are the main effects considered in the model. The levels “A2780” (*Data set*), “AlexNet” (*Model*), “4” (*Batch*), and “0” (*Augmentation*) were used as reference factor levels to calculate the intercept term of the linear model,  $\mu_0$ , which corresponds to the mean pIC<sub>50</sub> value for this combination of factor levels. The coefficients (i.e., slopes) for the other combinations of factor levels correspond to the difference between their mean pIC<sub>50</sub> value and the intercept. The error term,  $\varepsilon_{i,j,k,l,m}$ , corresponds to the random error of each pIC<sub>50</sub> value, defined as  $\varepsilon_{i,j,k,l,m} = pIC_{50i,j,k,l,m} - \text{mean}(pIC_{50i,j,k,l})$ . These errors are assumed to (1) be mutually independent, (2) have zero expectation value, and (3) have constant variance.

We trained ten models for each combination of factor levels, each time randomly assigning different sets of data points to the training, validation and test sets. The normality and homoscedasticity assumptions of the linear models were respectively assessed with (1) quantile–quantile (Q–Q) plots and (2) by plotting the fitted values against the residuals [77]. Homoscedasticity means that the residuals are equally dispersed across the range of the dependent variable used in the linear model. A systematic bias of the residuals would indicate that the errors are not random and that they contain predictive information that should be included in the model [78, 79].

To compare the performance of (1) the most predictive ConvNet for each data set and replication, (2) RF and (3) DNN models trained on Morgan fingerprints, and (4) the Ensemble models generated by averaging the predictions of the RF and ConvNet models, we also used a linear

model with two factors, namely *Data set* and *Model*. In this case, we only considered the results of the ConvNet architecture leading to the lowest RMSE value on the test set for each data set and replication.

$$\begin{aligned}
 pIC_{50} = & \text{Data set}_i + \text{Model}_j \\
 & + (\text{Data set} * \text{Model})_{ij} + \mu_0 + \varepsilon_{i,j,k} \quad (2) \\
 (i \in & \{1, \dots, N_{\text{datasets}} = 33\}; \\
 j \in & \{1, \dots, N_{\text{models}} = 4\})
 \end{aligned}$$

## Results and discussion

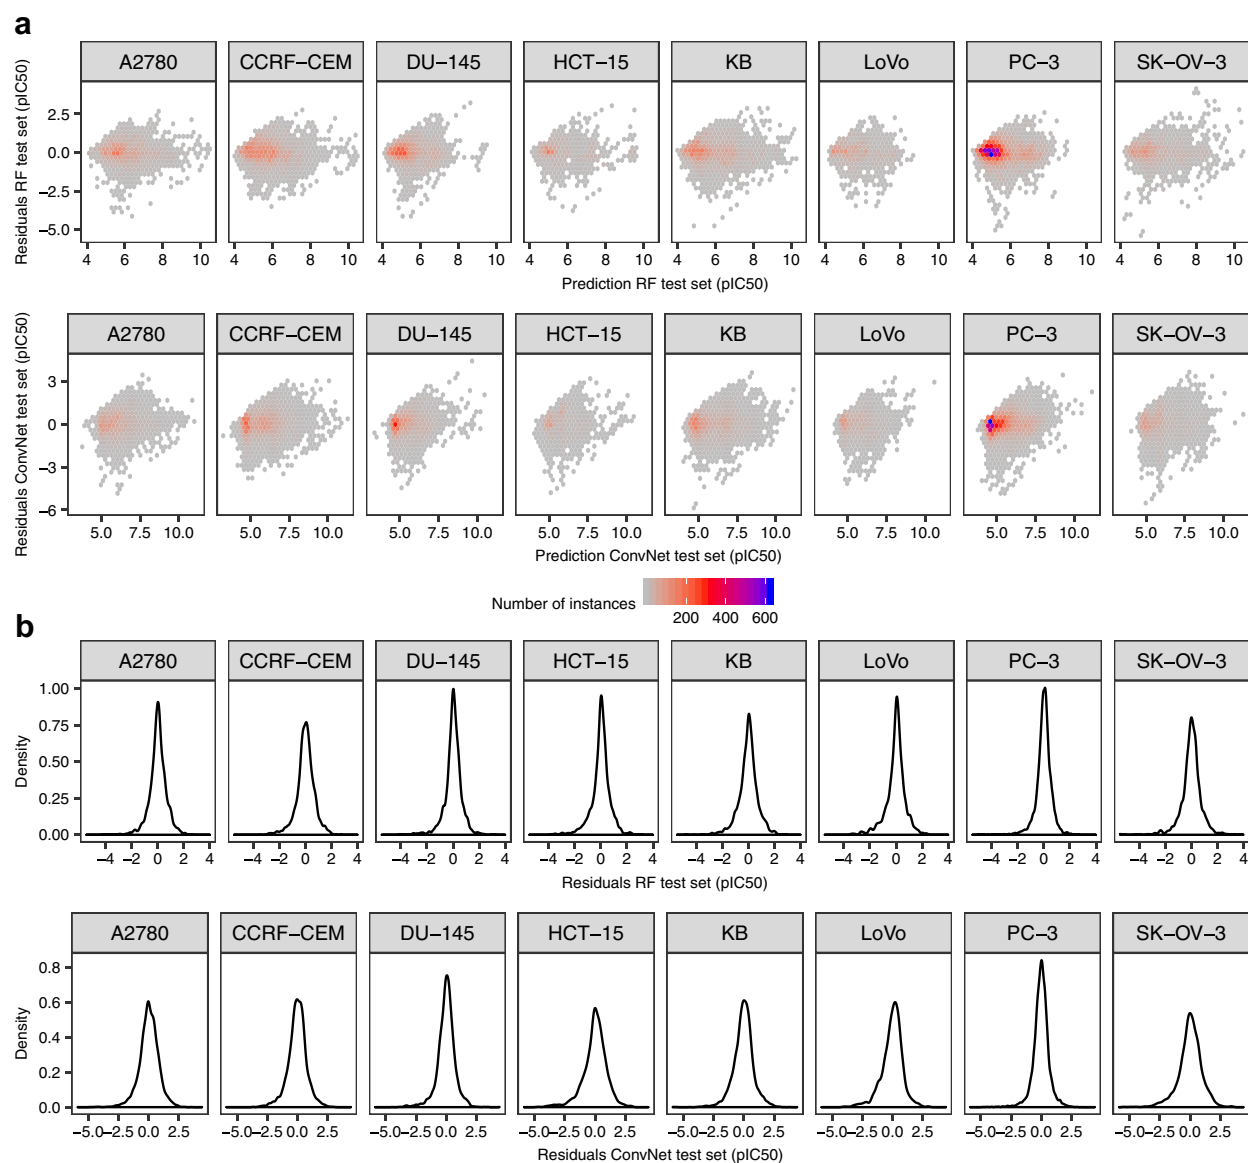
We initially evaluated the performance of ConvNets to predict the activity of compounds from their Kekulé structure representations using 8 cytotoxicity data sets. To this aim, we modelled these data sets using four widely-used architectures, namely AlexNet, DenseNet 201, ResNet-152, and VGG-19 with batch normalization (VGG-19-bn), and the extended versions thereof that we implemented by including four additional fully-connected layers after the convolutional layers (see Methods and Fig. 1). We obtained high performance on the test set for all networks, with mean RMSE values in the 0.65–0.96 pIC<sub>50</sub> range (Fig. 2). These errors in prediction are comparable to the uncertainty of heterogeneous IC<sub>50</sub> measurements in ChEMBL [8], and to the performance of drug sensitivity prediction models previously reported [15, 18, 80]. Notably, high performance was also obtained for data sets containing few hundred compounds (e.g., LoVo or HCT-15), suggesting that the framework proposed here is applicable to model small data sets.

In order to study the relative performance of the network architectures in a robust manner, we implemented a factorial design that we evaluated using a linear model (Eq. 1). The linear model displayed an R<sup>2</sup> value adjusted for the number of parameters of 0.68 ( $P < 10^{-12}$ ), thus indicating that the variables considered in our factorial design explain a large proportion of the variation observed in model performance, and hence, its coefficients provide valuable information to study the relative performance of the modelling strategies explored here in a statistically sound manner. Analysis of the model coefficients revealed that the performance of the extended versions of the architectures constantly led to a decrease in the RMSE values of ~ 5–10% ( $P < 10^{-12}$ ; Fig. 2), with ResNet-152, and VGG-19-bn constantly leading to the highest predictive models. Together, these results thus suggest that the four additional fully-connected layers we included in the architectures and the use of dropout regularization help palliate overfitting (Fig. 2), and hence, increase the generalization capabilities of the networks.

To ensure that the low RMSE values observed are not the consequence of simply predicting the mean value of the response variable, we examined the distributions of the residuals for the ConvNet and RF models (Fig. 3). These complementary analyses are important because, as we have previously shown for protein-ligand data sets [74], networks that fail to converge often simply predict the mean value of the dependent variable. Overall, we observed similar patterns for both modelling approaches (as shown in Fig. 3), with residuals centered around zero and generally showing homoscedasticity, i.e., displaying comparable variance across the entire bioactivity range. Examination of the residuals is also important when modelling imbalanced data sets, which is generally the case for data sets extracted from ChEMBL, because a large fraction of instances are annotated with pIC<sub>50</sub> values in the low micromolar range (4–5 pIC<sub>50</sub> units), and by simply predicting the mean value of the response variable one might already obtain low RMSE values (~ 1 pIC<sub>50</sub> units for these data sets, see yellow bars in Fig. 4). In such cases, the residuals would be heteroscedastic, displaying increasingly higher variances towards the low-nanomolar range (i.e., pIC<sub>50</sub> values of 8–9), which however was not the case for the models generated here. Together, these results thus indicate that compound images convey sufficient chemical information to model compound bioactivities across a wide dynamic range of pIC<sub>50</sub> values.

In addition, we performed Y-scrambling experiments using the 8 cytotoxicity data sets to ensure that the predictive power obtained by the ConvNets did not arise by chance. With this aim in mind, the bioactivity values for the training and validation set instances were shuffled before training. We observed R<sup>2</sup> values around 0 ( $P < 0.001$ ) for the observed against the predicted values on the test set for all the Y-scrambling experiments we performed. Therefore, these results indicate that the features extracted by the convolutional layers capture chemical information related to bioactivity, and that the high predictive power of the ConvNets is not a consequence of spurious correlations.

We previously showed that data augmentation represents a versatile approach to increase the predictive power of RF models trained on compound fingerprints [81]. Similarly, we here find a significant increase in performance for ConvNets trained on augmented data sets ( $P = 0.02$ ). In fact, the utilization of data augmentation during training led to the most predictive models in 68% of the cases; when considering the most predictive network for each data set and run only, we find that data augmentation was used in 91% of the cases.



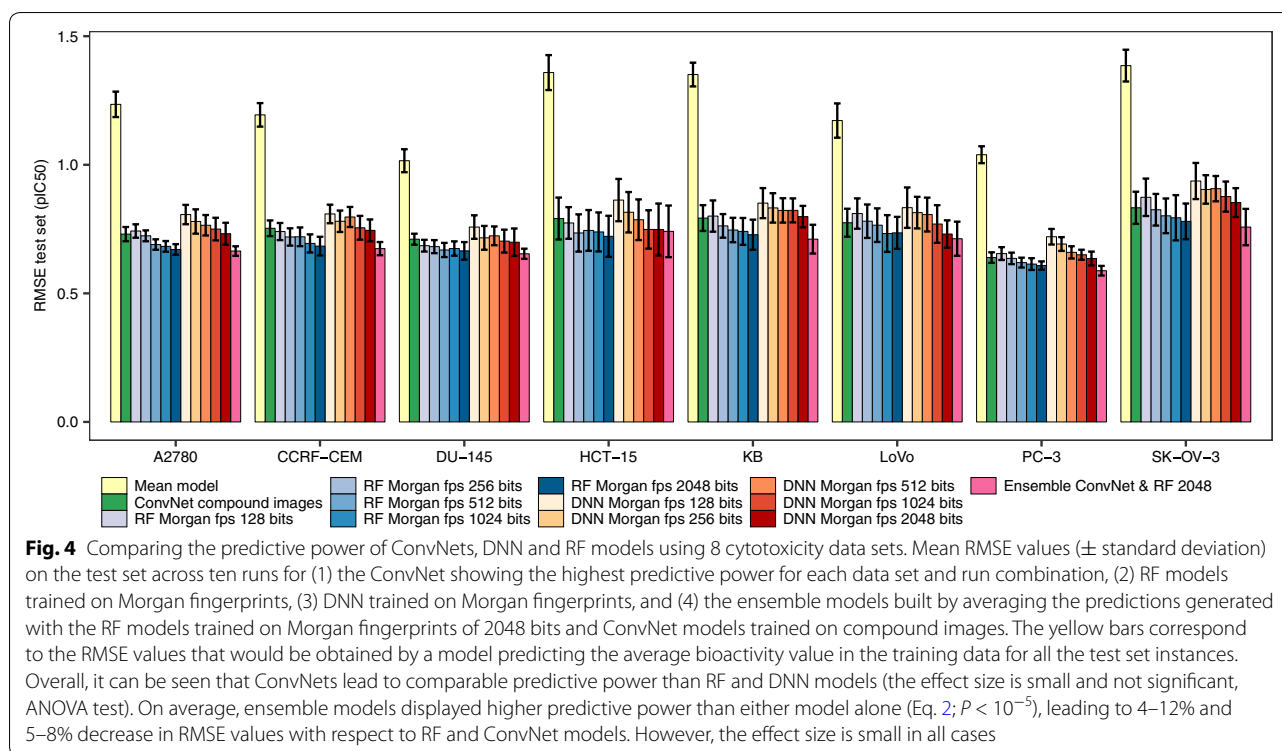
**Fig. 3** Analysis of the residuals. Residuals for the ConvNets (top panels) and RF (bottom panels) models for the cytotoxicity data sets. Overall, the residuals for both types of models show comparable variance across the bioactivity range (a) and are centered around zero (b), indicating that compound images permit to model the activity of small molecules across a dynamic range of pIC<sub>50</sub> values

Overall, these results indicate that the extraction of chemical information by the ConvNets is robust against rotations of the compound images, and that data augmentation helps improve chemical-structure activity modelling based on compound images [81].

Next, we compared the predictive power of the ConvNets to that of RF and DNN models trained on Morgan fingerprints of increasingly higher dimensionality (from 128 to 2048 bits) using the factorial design described in Eq. 2. The linear model in this case showed an adjusted  $R^2$  value of 0.97, suggesting that

the covariates we considered account for most of the variability in model performance. Overall, we did not find significant differences in performance between RF models, DNN trained on circular fingerprints and ConvNets trained on compound images ( $P = 0.76$ ; Figs. 4, 5). The former models are using Morgan FP and RF or DNN, which have previously been shown to generate models with high predictive power in benchmarking studies of compound descriptors and algorithms [81–83]. Taken together, these results suggest that compound images provide sufficient predictive signal to

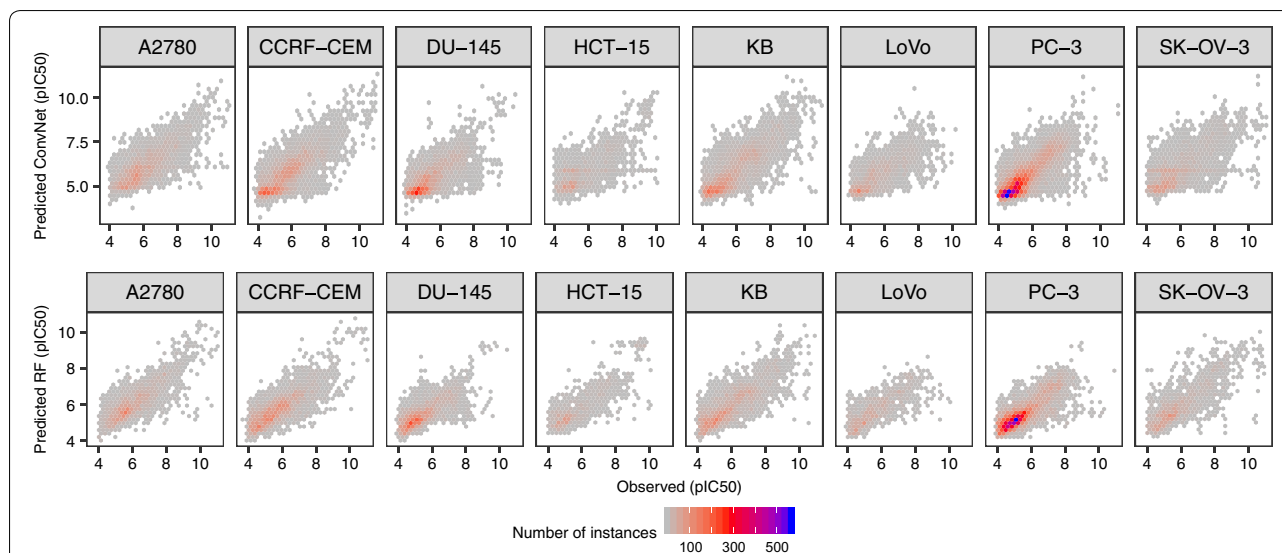




generate ConvNets with comparable predictive power to state-of-the-art methods, even for small data sets of few hundred compounds.

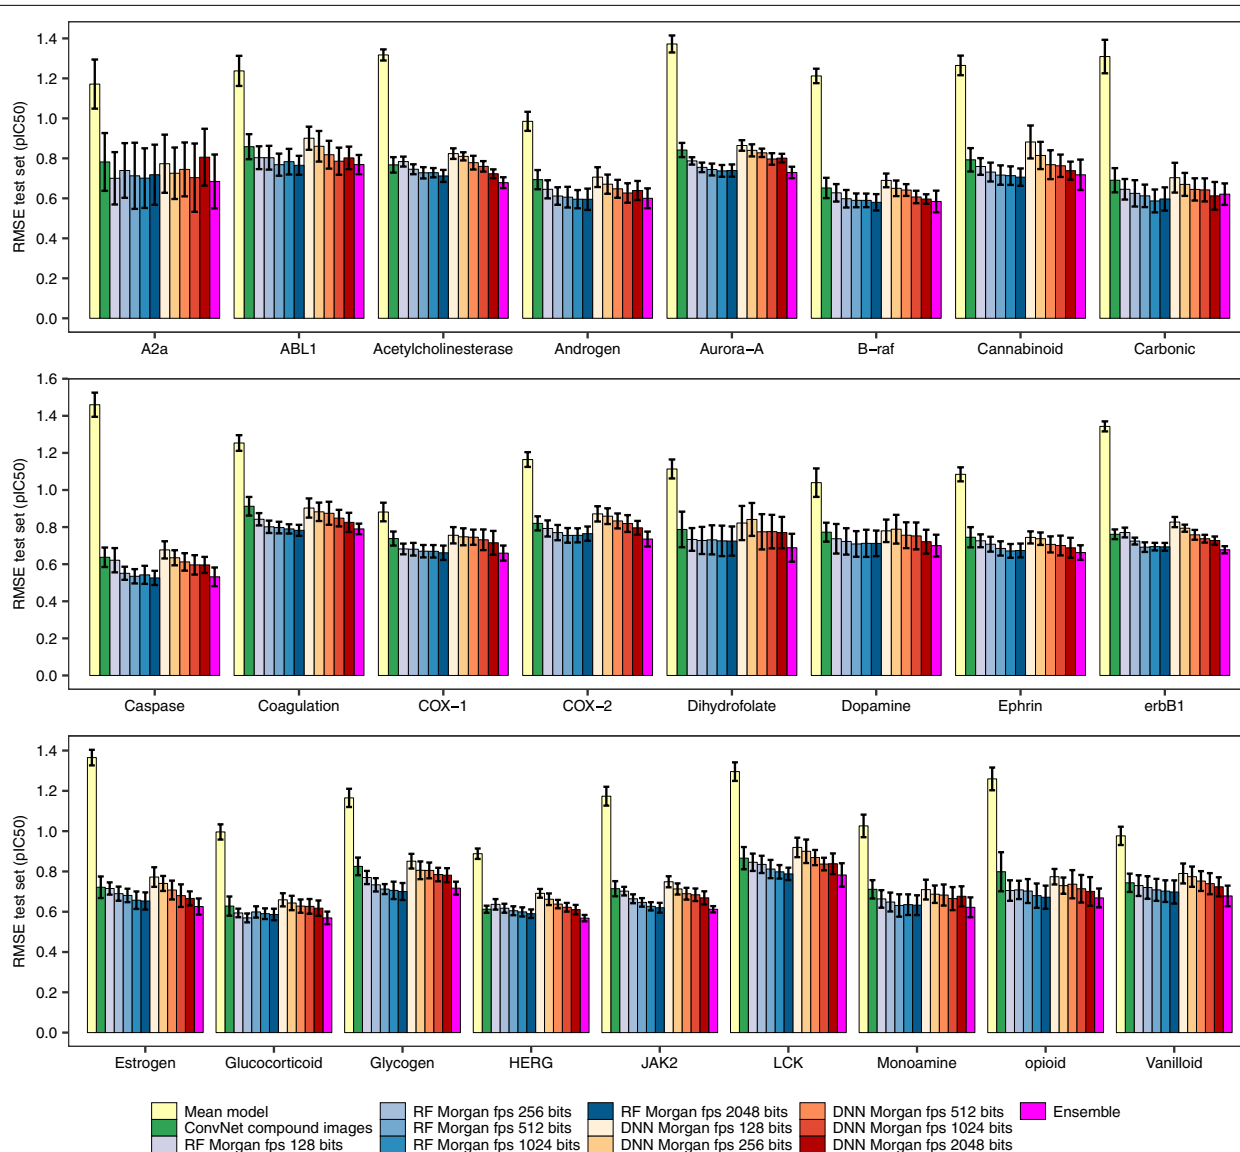
As an additional validation of our modelling framework, we extended our analysis to 25 protein

target data sets (Table 2). We trained ConvNets using the ResNet-152 and VGG-19-bn architectures given their higher performance when modelling the cytotoxicity data sets described above. Overall, we obtained comparable performance for ConvNets, RF and DDN models (Fig. 6),



with effect sizes across algorithms in the 0.03–0.09 RMSE (i.e.,  $pIC_{50}$ ) units range. Y-scrambling experiments for these data sets also led to  $R^2$  values around 0 ( $P < 0.001$ ). We next capitalized on the large number of compounds annotated with bioactivity data for both COX-1 and COX-2 [84, 85] to model COX isoform selectivity using multi-task ConvNets trained on compound images. Multi-task ConvNets displayed comparable performance

to single-task ConvNets trained using either COX-1 or COX-2 data, with RMSE on the test set in the 0.72–0.75 range (Table 4), which are comparable to the uncertainty in heterogeneous  $pIC_{50}$  data extracted from ChEMBL [86]. Together, these results indicate that ConvNets extract structural aspects related to compound activity from compound images, which in turn enable the modeling of diverse bioactivity read-outs (compound potency



**Fig. 6** Benchmarking the predictive power of ConvNets, DNN and RF models on 25 protein target data sets. Mean RMSE values ( $\pm$  standard deviation) on the test set across ten runs for (1) the ConvNet showing the highest predictive power for each data set and run combination, (2) RF models trained on Morgan fingerprints, (3) DNN trained on Morgan fingerprints, and (4) the ensemble models built by averaging the predictions generated with the RF models trained on Morgan fingerprints of 2048 bits and ConvNet models trained on compound images. As in Fig. 4, the yellow bars correspond to the RMSE values that would be obtained by a model predicting the average bioactivity value in the training data for all the test set instances. Similar to the results obtained for the cytotoxicity data sets, we obtained ConvNets with comparable predictive power to RF and DNN models. As in the case of the cytotoxicity data sets, ensemble models often displayed higher predictive power than either model alone (Eq. 2;  $P < 10^{-5}$ )

and cell growth inhibition), measured in target systems of increasing complexity, from purified proteins to cell cultures.

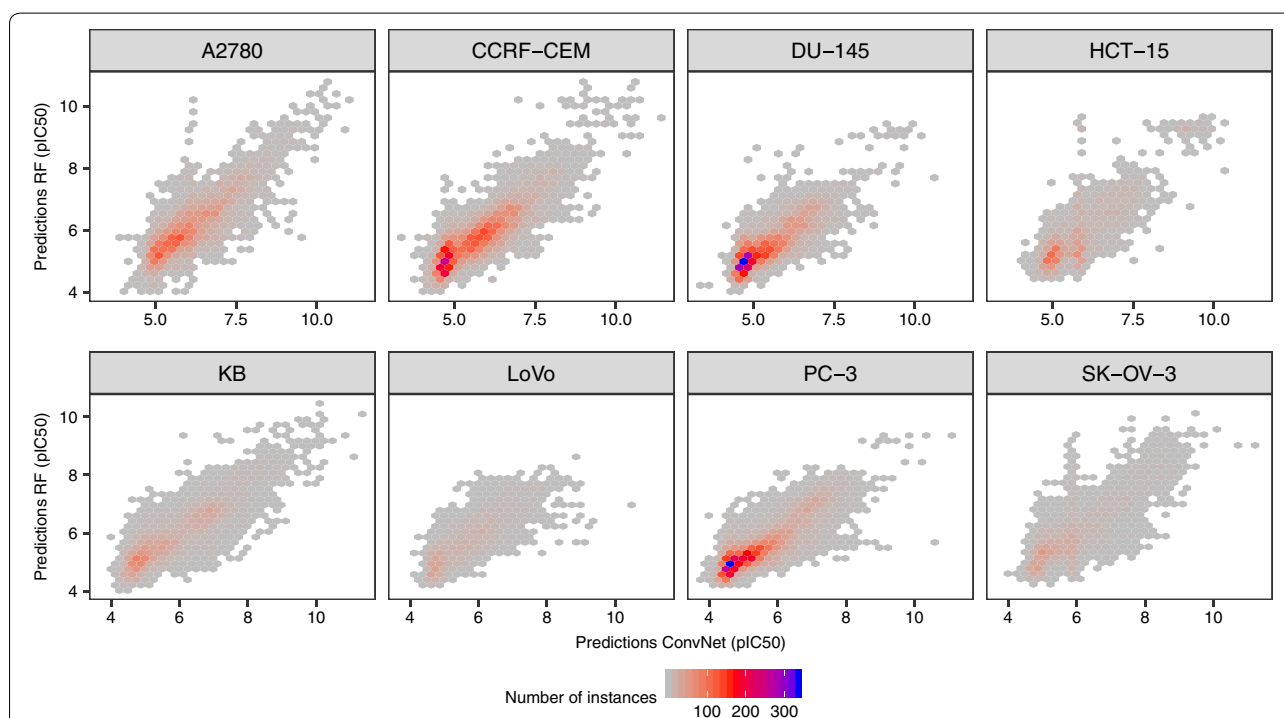
To further characterize the differences between RF and ConvNets, we firstly assessed the correlation between the predicted values calculated for the same test set instances using models trained on the same data splits. We found (as shown in Fig. 7) that the predictions of both models are highly correlated for all data sets, with  $R^2$  values in the 0.80–0.89 range (Pearson's correlation coefficient;  $P < 0.05$ ), thus indicating that the predictions calculated with the RF models explain a large fraction of the variance observed for the predictions calculated with the ConvNets, and *vice versa*. Analysis of the correlation of the absolute error in prediction for each test set instance

however revealed that the error profiles of RF and ConvNets are only moderately correlated ( $R^2$  in the 0.58–0.65 range,  $P < 0.05$ ; Fig. 8). From the latter, we hypothesized that combining the predictions generated by each modelling technique into a model ensemble might lead to increased predictive power [84]. In fact, ensemble models built by averaging the predictions generated by RF and ConvNet models displayed higher predictive power in some cases, leading to 4–12% and 5–8% decrease in RMSE values with respect to RF and ConvNet models, respectively ( $P < 10^{-5}$ ; pink bars in Figs. 4 and 6). In contrast to previous analyses [51], where compound fingerprints and related representations were often thought to contain most information related to bioactivity [87], our results indicate that Morgan fingerprints and the features extracted from compound images with the ConvNets convey complementary predictive signal for some data sets, thus permitting to obtain more accurate predictions than either model alone by combining them into a model ensemble.

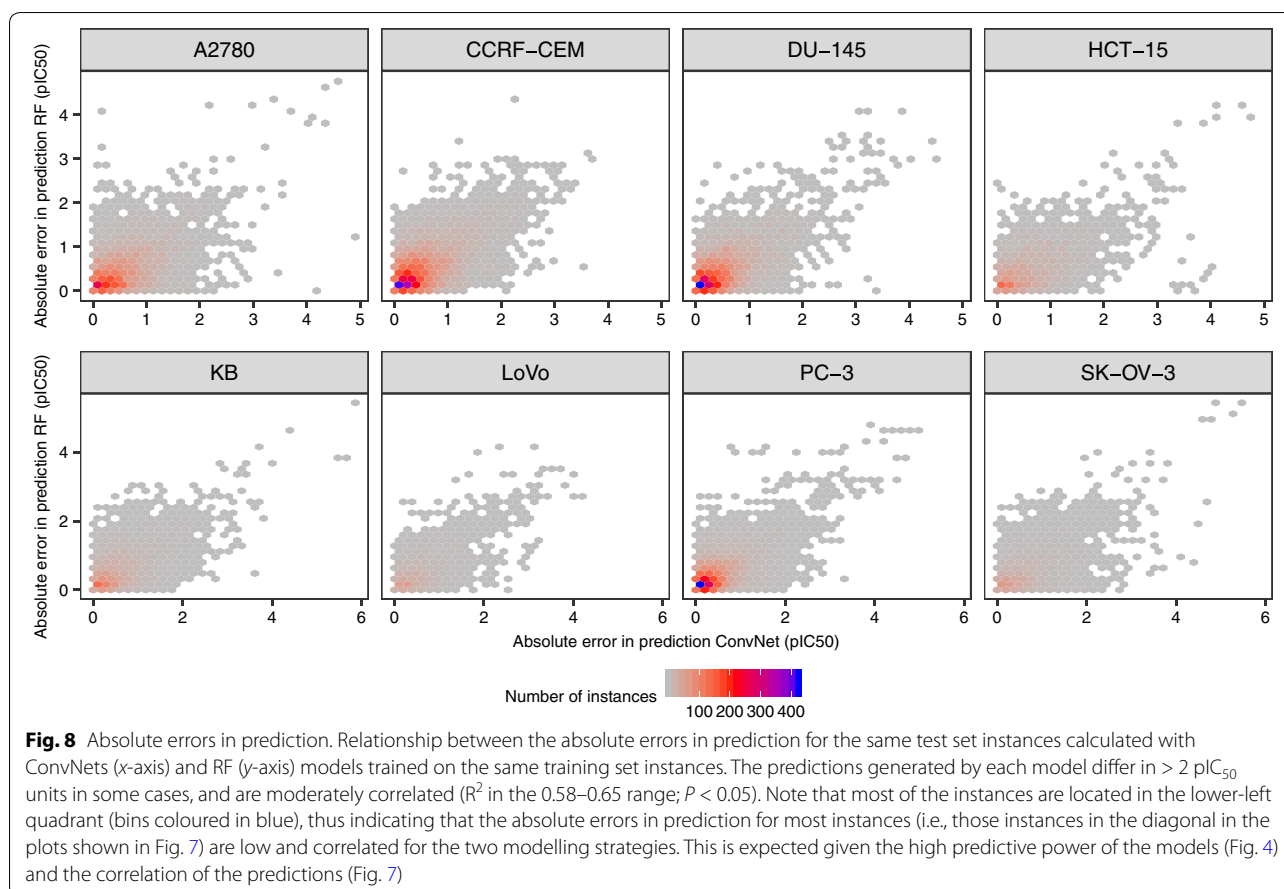
In this work, we show using 33 diverse data sets extracted from ChEMBL database that a proper design and parametrization of ConvNets is sufficient to generate highly predictive models trained on images of structural compound representations sketched using standard functionalities of commonly used software packages (e.g.,

**Table 4 Predictive power on the test set of multi-task and single-task models trained on the COX-1 and COX-2 data sets**

Model	RMSE $\pm$ SD on COX-1	RMSE $\pm$ SD on COX-2
Multi-target COX-1 and COX-2	0.73 $\pm$ 0.05	0.75 $\pm$ 0.05
Single-target COX-1	0.73 $\pm$ 0.06	NA
Single-target COX-2	NA	0.72 $\pm$ 0.05



**Fig. 7** RF and ConvNet predictions on the test set. Correlation between the predictions for the test set compounds calculated with RF models and ConvNets trained on the same training set instances. Overall, the predictions show a positive and significant correlation ( $P < 0.05$ ; Pearson's correlation coefficient values in the 0.72–0.84 range). The predictions for the ten runs are shown



RDkit). Therefore, exploiting such networks, which were designed for general image recognition tasks, and pre-trained on unrelated image data sets, represents a versatile approach to model compound activity directly from Kekulé structure representations in a purely data-driven fashion. However, it is paramount to note that the computational footprint of ConvNets still represents a major limitation of this approach: whereas training the RF models for these data sets required 6–14 s *per model* using 16 CPU cores and no parameter optimization, training times *per epoch* for the ConvNets were in the 15–64 s range (i.e., 150–640 min *per model* using one GPU card and 16 CPU cores for image processing).

While the computation of compound descriptors has traditionally relied on predefined rules or prior knowledge of chemical properties, bioactivity profiles or topological information of compounds, among others [88–91], the descriptors calculated by the convolutional layers of ConvNets represent an automatic and data-driven approach to derive features directly from chemical structure representations [42], as we do here, or from image representations of a predefined set of molecular and topological features [41, 42]. As we show in this

study, these compound features permit to model compound bioactivity with high accuracy even on a continuous scale. However, image-derived features are generally harder to interpret than more traditional descriptors, e.g., keyed Morgan fingerprints [84], although few methods to interpret convolutional graphs have been previously proposed [41, 92]. We anticipate that extending the work presented here by including 3D representations of compounds and binding sites using 3D convolutional neural networks to account for conformational changes of small molecules and protein dynamics, respectively, will likely improve compound activity modelling [93–97].

Previous work using compound images and neural networks to model compound toxicity has shown that using a molecular representation where atoms are colored yields high predictive power [23]. We note that there are countless sketching protocols to represent molecules, and hence, future benchmarking studies will be needed to thoroughly examine their predictive signal. Similarly, elucidating the most convenient strategies to perform data augmentation is an area of intense research [98–100], also for chemical structure-activity modelling [81, 101]. In the case of ConvNets, multiple representations

of the same molecules generated using diverse sketching schemes (e.g., using diverse SMILES encoding rules [101]) might be implemented to perform data augmentation. Therefore, future comparative studies of data augmentation strategies will also be needed to determine the most appropriate one for bioactivity modelling using ConvNets.

The neural network architectures used in this study require the input images to be of size 224x224, as modelling larger images would result in a computationally intractable increase in the number of parameters. Therefore, we generated images of that size for all compounds. Such an approach however results in larger representations for small molecules as compared to larger ones: the same chemical moiety might span a larger or smaller region in the images depending on the size of the molecule in which it appears. To account for this issue, images could be cropped to enlarge functional groups as a data augmentation strategy during the learning process. In this study, we did not investigate this further as the generalization capability of the networks we generated was comparable to that of RF and fully-connected networks trained on Morgan fingerprints. Thus, the influence on model performance of the relative size of the representations of chemical moieties and functional groups across molecules remains to be thoroughly examined.

Finally, future work will also be required to evaluate whether ConvNets trained on both compound and cellular images lead to more accurate modelling of compound activity on cancer cell lines, as well as other output variables (i.e., toxicity), than current modelling approaches based on gene expression or mutation profiles [15, 16, 18, 102].

## Conclusions

In this contribution, we introduce *KekuleScope*, a framework to model compound bioactivity on a continuous scale using extended versions of four widely-used architectures trained on Kekulé structure representations without requiring any image preprocessing or network engineering steps. The generated models achieve comparable performance to RF and DNN models trained on circular fingerprints, and to the estimated experimental uncertainty of the input data. Our work shows that Kekulé representations can be harnessed to derive robust models without requiring any additional descriptor calculation. In addition, we show that the chemical information extracted by the convolutional layers of the ConvNets is often complementary to that provided by Morgan fingerprints, which enables the generation of model ensembles with significantly higher predictive power than either RF models or ConvNets alone,

although the effect size is small. The framework proposed here is generally applicable across endpoints, and it is expected that also on other datasets the combination of models will lead to increases in performance.

## Acknowledgements

None.

## Authors' contributions

IC-C conceived and designed the study, IC-C implemented the models, interpreted and analyzed the results. IC-C generated the figures and wrote the paper with substantial input from AB. Both authors read and approved the final manuscript.

## Funding

This Project has received funding from the European Union's Framework Programme For Research and Innovation Horizon 2020 (2014–2020) under the Marie Curie Skłodowska-Curie Grant Agreement No. 703543 (I.C.C.).

## Availability of data and materials

The code and data sets used in this study is available free of charge at <https://github.com/isidro/kekulescope>.

## Ethics approval and consent to participate

Not applicable, as no human participants, data or tissue or animals are involved in this study.

## Consent for publication

The authors agree to the conditions of submission, BioMed Central's copyright and license agreement and article-processing charge (APC).

## Competing interests

The authors declare that they have no competing interests.

Received: 19 February 2019 Accepted: 9 June 2019

Published online: 19 June 2019

## References

1. Basu A, Bodycombe NE, Cheah JH et al (2013) An interactive resource to identify cancer genetic and lineage dependencies targeted by small molecules. *Cell* 154:1151–1161. <https://doi.org/10.1016/j.cell.2013.08.003>
2. Seashore-Ludlow B, Rees MG, Cheah JH et al (2015) Harnessing connectivity in a large-scale small-molecule sensitivity dataset. *Cancer Discov* 5:1210–1223. <https://doi.org/10.1158/2159-8290.CD-15-0235>
3. Shoemaker RH (2006) The NCI60 human tumour cell line anticancer drug screen. *Nat Rev Cancer* 6:813–823. <https://doi.org/10.1038/nrc1951>
4. Yang W, Soares J, Greninger P et al (2013) Genomics of drug sensitivity in cancer (GDSC): a resource for therapeutic biomarker discovery in cancer cells. *Nucleic Acids Res* 41:D955–D961. <https://doi.org/10.1093/nar/gks1111>
5. Barretina J, Caponigro G, Stransky N et al (2012) The cancer cell line encyclopedia enables predictive modelling of anticancer drug sensitivity. *Nature* 483:603–607. <https://doi.org/10.1038/nature11003>
6. Garnett MJ, McDermott U (2014) The evolving role of cancer cell line-based screens to define the impact of cancer genomes on drug response. *Curr Opin Genet Dev* 24:114–119. <https://doi.org/10.1016/j.gde.2013.12.002>
7. Knouse KA, Lopez KE, Bachofner M, Amon A (2018) Chromosome segregation fidelity in epithelia requires tissue architecture. *Cell* 175:200–211. e13. <https://doi.org/10.1016/j.cell.2018.07.042>
8. Cortés-Ciriano I, Bender A (2015) How consistent are publicly reported cytotoxicity data? Large-scale statistical analysis of the concordance of public independent cytotoxicity measurements. *ChemMedChem* 11:57–71. <https://doi.org/10.1002/cmdc.201500424>

9. Ben-David U, Siranosian B, Ha G et al (2018) Genetic and transcriptional evolution alters cancer cell line drug response. *Nature* 560:325–330. <https://doi.org/10.1038/s41586-018-0409-3>
10. Iorio F, Knijnenburg TA, Vis DJ et al (2016) A landscape of pharmacogenomic interactions in cancer. *Cell* 168:740–754. <https://doi.org/10.1016/j.cell.2016.06.017>
11. Najgebauer H, Yang M, Francies H et al (2018) CELLector: genomics guided selection of cancer in vitro models. *bioRxiv*. <https://doi.org/10.1101/275032>
12. Gorthi A, Romero JC, Loranc E et al (2018) EWS–FLI1 increases transcription to cause R-loops and block BRCA1 repair in Ewing sarcoma. *Nature* 555:387–391. <https://doi.org/10.1038/nature25748>
13. Menden MP, Casale FP, Stephan J et al (2018) The germline genetic component of drug sensitivity in cancer cell lines. *Nat Commun* 9:3385. <https://doi.org/10.1038/s41467-018-05811-3>
14. Rees MG, Seashore-Ludlow B, Cheah JH et al (2016) Correlating chemical sensitivity and basal gene expression reveals mechanism of action. *Nat Chem Biol* 12:109–116. <https://doi.org/10.1038/nchembio.1986>
15. Cortés-Ciriano I, van Westen GJP, Bouvier G et al (2016) Improved large-scale prediction of growth inhibition patterns using the NCI60 cancer cell line panel. *Bioinformatics* 32:85–95. <https://doi.org/10.1093/bioinformatics/btv529>
16. Cortes-Ciriano I, Mervin L, Bender A (2017) Current trends in drug sensitivity prediction. *Curr Pharm Des* 22:6918–6927. <https://doi.org/10.2174/1381612822666161026154430>
17. Altman RB (2015) Predicting cancer drug response: advancing the DREAM. *Cancer Discov* 5:237–238. <https://doi.org/10.1158/2159-8290.CD-15-0093>
18. Menden MP, Iorio F, Garnett M et al (2013) Machine learning prediction of cancer cell sensitivity to drugs based on genomic and chemical properties. *PLoS ONE* 8:e61318. <https://doi.org/10.1371/journal.pone.0061318>
19. Geleher P, Cox NJ, Huang RS (2014) Clinical drug response can be predicted using baseline gene expression levels and in vitro drug sensitivity in cell lines. *Genome Biol* 15:R47. <https://doi.org/10.1186/gb-2014-15-3-r47>
20. Naulaerts S, Dang CC, Ballester PJ (2017) Precision and recall oncology: combining multiple gene mutations for improved identification of drug-sensitive tumours. *Oncotarget* 8:97025–97040. <https://doi.org/10.18632/oncotarget.20923>
21. Chen H, Engkvist O, Wang Y et al (2018) The rise of deep learning in drug discovery. *Drug Discov Today* 23:1241–1250. <https://doi.org/10.1016/j.drudis.2018.01.039>
22. Mayr A, Klambauer G, Unterthiner T, Hochreiter S (2016) DeepTox: toxicity prediction using deep learning. *Front Environ Sci* 3:80. <https://doi.org/10.3389/fenvs.2015.00080>
23. Fernandez M, Ban F, Woo G et al (2018) Toxic colors: the use of deep learning for predicting toxicity of compounds merely from their graphic images. *J Chem Inf Model* 58:1533–1543. <https://doi.org/10.1021/acs.jcim.8b00338>
24. Ramsundar B, Kearnes S, Riley P et al (2015) Massively multitask networks for drug discovery. *arXiv:1502.02072*. Accessed 20 July 2018
25. Dahl GE, Jaitly N, Salakhutdinov R (2014) Multi-task neural networks for QSAR predictions. *arXiv:1406.1231*. Accessed 19 July 2018
26. Preuer K, Lewis RPI, Hochreiter S et al (2018) DeepSynergy: predicting anti-cancer drug synergy with deep learning. *Bioinformatics* 34:1538–1546. <https://doi.org/10.1093/bioinformatics/btx806>
27. Koutsoukas A, Monaghan KJ, Li X, Huan J (2017) Deep-learning: investigating deep neural networks hyper-parameters and comparison of performance to shallow methods for modeling bioactivity data. *J Cheminform* 9:42. <https://doi.org/10.1186/s13321-017-0226-y>
28. Altae-Tran H, Ramsundar B, Pappu AS, Pande V (2017) Low data drug discovery with one-shot learning. *ACS Cent Sci* 3:283–293. <https://doi.org/10.1021/acscentsci.6b00367>
29. Subramanian G, Ramsundar B, Pande V, Denny RA (2016) Computational modeling of  $\beta$ -secretase 1 (BACE-1) inhibitors using ligand based approaches. *J Chem Inf Model* 56:1936–1949. <https://doi.org/10.1021/acs.jcim.6b00290>
30. Korotcov A, Tkachenko V, Russo DP, Ekins S (2017) Comparison of deep learning with multiple machine learning methods and metrics using diverse drug discovery data sets. *Mol Pharm* 14:4462–4475. <https://doi.org/10.1021/acs.molpharmaceut.7b00578>
31. Putin E, Asadulaev A, Vanhaelen Q et al (2018) Adversarial threshold neural computer for molecular de novo design. *Mol Pharm* 15:4386–4397. <https://doi.org/10.1021/acs.molpharmaceut.7b01137>
32. Olivecrona M, Blaschke T, Engkvist O, Chen H (2017) Molecular de-novo design through deep reinforcement learning. *J Cheminform* 9:48. <https://doi.org/10.1186/s13321-017-0235-x>
33. Gómez-Bombarelli R, Wei JN, Duvenaud D et al (2018) Automatic chemical design using a data-driven continuous representation of molecules. *ACS Cent Sci* 4:268–276. <https://doi.org/10.1021/acscentsci.7b00572>
34. Segler MHS, Kogej T, Tyrchan C, Waller MP (2018) Generating focused molecule libraries for drug discovery with recurrent neural networks. *ACS Cent Sci* 4:120–131. <https://doi.org/10.1021/acscentsci.7b00512>
35. Ma J, Sheridan RP, Liaw A et al (2015) Deep neural nets as a method for quantitative structure–activity relationships. *J Chem Inf Model* 55:263–274. <https://doi.org/10.1021/ci500747n>
36. Lecun Y, Bengio Y, Hinton G (2015) Deep learning. *Nature* 521:436–444. <https://doi.org/10.1038/nature14539>
37. Lecun Y, Bottou L, Bengio Y, Haffner P (1998) Gradient-based learning applied to document recognition. *Proc IEEE* 86:2278–2324. <https://doi.org/10.1109/5.726791>
38. Le Cun Y, Boser B, Denker JS et al (1989) Handwritten digit recognition with a back-propagation network. In: *Proceedings of the 2nd International Conference on Neural Information Processing Systems*, pp 396–404
39. Wainberg M, Merico D, DeLong A, Frey BJ (2018) Deep learning in biomedicine. *Nat Biotechnol* 36:829–838. <https://doi.org/10.1038/nbt.4233>
40. Liu K, Sun X, Jia L et al (2018) Chemi-net: a graph convolutional network for accurate drug property prediction. *arXiv:1803.06236*. Accessed 8 July 2018
41. Duvenaud D, Maclaurin D, Aguilera-Iparraguirre J et al (2015) Convolutional networks on graphs for learning molecular fingerprints. *arXiv:1509.09292*. Accessed 8 July 2018
42. Kearnes S, McCloskey K, Berndl M et al (2016) Molecular graph convolutions: moving beyond fingerprints. *J Comput Aided Mol Des* 30:595–608. <https://doi.org/10.1007/s10822-016-9938-8>
43. Cooper LA, Demicco EG, Saltz JH et al (2018) PanCancer insights from The Cancer Genome Atlas: the pathologist's perspective. *J Pathol* 244:512–524. <https://doi.org/10.1002/path.5028>
44. Coudray N, Ocampo PS, Sakellaropoulos T et al (2018) Classification and mutation prediction from non-small cell lung cancer histopathology images using deep learning. *Nat Med*. <https://doi.org/10.1038/s41591-018-0177-5>
45. Yu K-H, Zhang C, Berry GJ et al (2016) Predicting non-small cell lung cancer prognosis by fully automated microscopic pathology image features. *Nat Commun* 7:12474. <https://doi.org/10.1038/ncomms12474>
46. Yu K-H, Beam AL, Kohane IS (2018) Artificial intelligence in healthcare. *Nat Biomed Eng* 2:719–731. <https://doi.org/10.1038/s41551-018-0305-z>
47. Poplin R, Chang P-C, Alexander D et al (2018) A universal SNP and small-indel variant caller using deep neural networks. *Nat Biotechnol* 36:983. <https://doi.org/10.1038/nbt.4235>
48. Scheeder C, Heigwer F, Boutros M (2018) Machine learning and image-based profiling in drug discovery. *Curr Opin Syst Biol* 10:43–52. <https://doi.org/10.1016/j.coisb.2018.05.004>
49. Kraus OZ, Ba JL, Frey BJ (2016) Classifying and segmenting microscopy images with deep multiple instance learning. *Bioinformatics* 32:i52–i59. <https://doi.org/10.1093/bioinformatics/btw252>
50. Hofmarcher M, Rumetshofer E, Clevert D-AA et al (2019) Accurate prediction of biological assays with high-throughput microscopy images and convolutional networks. *J Chem Inf Model* 59:1163–1171. <https://doi.org/10.1021/acs.jcim.8b00670>
51. Goh GB, Siegel C, Vishnu A, Hodas NO (2017) Using rule-based labels for weak supervised learning: a ChemNet for transferable chemical property prediction. [https://doi.org/10.475/123\\_4](https://doi.org/10.475/123_4). *arXiv:1712.02734*. Accessed 8 July 2018
52. Goh GB, Siegel C, Vishnu A et al (2018) How much chemistry does a deep neural network need to know to make accurate predictions? In:

- Proceedings—2018 IEEE winter conference on applications of computer vision, WACV 2018, pp 1340–1349
53. Simm J, Klambauer G, Arany A et al (2018) Repurposing high-throughput image assays enables biological activity prediction for drug discovery. *Cell Chem Biol* 25:611–618.e3. <https://doi.org/10.1016/j.chembiol.2018.01.015>
  54. Goh GB, Siegel C, Vishnu A et al (2017) Chemception: a deep neural network with minimal chemistry knowledge matches the performance of expert-developed QSAR/QSPR models. [arXiv:1706.06689](https://arxiv.org/abs/1706.06689). Accessed 8 July 2018
  55. Szegedy C, Ioffe S, Vanhoucke V, Alemi A (2016) Inception-v4, inception-ResNet and the impact of residual connections on learning. [arXiv:1602.07261](https://arxiv.org/abs/1602.07261). Accessed 8 July 2018
  56. Wu Z, Ramsundar B, Feinberg EN et al (2017) MoleculeNet: a benchmark for molecular machine learning. [arXiv:1703.00564](https://arxiv.org/abs/1703.00564). Accessed 8 July 2018
  57. Krizhevsky A, Sutskever I, Hinton GE (2012) Imagenet classification with deep convolutional neural networks. *Adv Neural Inf Process Syst*. 1:1097–1105. <https://doi.org/10.1145/3065386>
  58. Huang G, Liu Z, van der Maaten L, Weinberger KQ (2016) Densely connected convolutional networks. [arXiv:1608.06993](https://arxiv.org/abs/1608.06993). Accessed 8 July 2018
  59. He K, Zhang X, Ren S, Sun J (2015) Deep residual learning for image recognition. [arXiv:1512.03385](https://arxiv.org/abs/1512.03385). Accessed 8 July 2018
  60. Simonyan K, Zisserman A (2014) Very deep convolutional networks for large-scale image recognition. ICLR. <https://doi.org/10.1016/j.infsf.2008.09.005>
  61. Rogers D, Hahn M (2010) Extended-connectivity fingerprints. *J Chem Inf Model* 50:742–754. <https://doi.org/10.1021/ci100050t>
  62. Morgan HL (1965) The generation of a unique machine description for chemical structures—a technique developed at chemical abstracts service. *J Chem Doc* 5:107–113. <https://doi.org/10.1021/c160017a018>
  63. Nowotka M, Papadatos G, Davies M et al (2016) Want drugs? Use python. [arXiv:1607.00378](https://arxiv.org/abs/1607.00378). Accessed 10 July 2018
  64. Davies M, Nowotka M, Papadatos G et al (2015) ChEMBL web services: streamlining access to drug discovery data and utilities. *Nucleic Acids Res* 43:W612–W620. <https://doi.org/10.1093/nar/gkv352>
  65. Gaulton A, Bellis LJ, Bento AP et al (2011) ChEMBL: a large-scale bioactivity database for drug discovery. *Nucleic Acids Res* 40:1100–1107. <https://doi.org/10.1093/nar/gkr777>
  66. Cherkasov A, Muratov EN, Fourches D et al (2014) QSAR modeling: where have you been? Where are you going to? *J Med Chem* 57:4977–5010. <https://doi.org/10.1021/jm4004285>
  67. O’Boyle NM, Sayle RA (2016) Comparing structural fingerprints using a literature-based similarity benchmark. *J Cheminform* 8:36. <https://doi.org/10.1186/s13321-016-0148-0>
  68. Fourches D, Muratov E, Tropsha A (2010) Trust, but verify: on the importance of chemical structure curation in cheminformatics and QSAR modeling research. *J Chem Inf Model* 50:1189–1204. <https://doi.org/10.1021/ci100176x>
  69. Landrum G (2017) RDKit: open-source cheminformatics. <https://www.rdkit.org/>. Accessed 12 Jan 2017
  70. Deng J, Dong W, Socher R et al (2009) ImageNet: a large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition, pp 248–255
  71. Paszke A, Chanan G, Lin Z et al (2017) Automatic differentiation in PyTorch. *Adv Neural Inf Process Syst* 30:1–4
  72. Sutskever I, Martens J, Dahl G, Hinton G (2013) On the importance of initialization and momentum in deep learning. In: Proceedings of the 30th international conference on machine learning, PMLR 28, pp 1139–1147
  73. Srivastava N, Hinton G, Krizhevsky A, Salakhutdinov R (2014) Dropout: a simple way to prevent neural networks from overfitting. *J Mach Learn Res* 15:1929–1958
  74. Cortés-Ciriano I, Bender A (2019) Deep confidence: a computationally efficient framework for calculating reliable prediction errors for deep neural networks. *J Chem Inf Model* 59:1269–1281. <https://doi.org/10.1021/acs.jcim.8b00542>
  75. Pedregosa F, Varoquaux G, Gramfort A et al (2011) Scikit-learn: machine learning in python. *J Mach Learn Res* 12:2825–2830
  76. Sheridan RP (2012) Three useful dimensions for domain applicability in QSAR models using random forest. *J Chem Inf Model* 52:814–823. <https://doi.org/10.1021/ci300004n>
  77. Winer B, Brown D, Michels K (1991) Statistical principles in experimental design, 3rd edn. McGraw-Hill, New York
  78. Roy K, Ambure P, Aher RB (2017) How important is to detect systematic error in predictions and understand statistical applicability domain of QSAR models? *Chemom Intell Lab Syst* 162:44–54. <https://doi.org/10.1016/j.chemolab.2017.01.010>
  79. Roy K, Das RN, Ambure P, Aher RB (2016) Be aware of error measures. Further studies on validation of predictive QSAR models. *Chemom Intell Lab Syst* 152:18–33. <https://doi.org/10.1016/J.CHEMO LAB.2016.01.008>
  80. Ammad-ud-din M, Georgii E, Gönen M et al (2014) Integrative and personalized QSAR analysis in cancer by kernelized Bayesian matrix factorization. *J Chem Inf Model* 54:2347–2359. <https://doi.org/10.1021/ci500152b>
  81. Cortes-Ciriano I, Bender A (2015) Improved chemical structure–activity modeling through data augmentation. *J Chem Inf Model* 55:2682–2692. <https://doi.org/10.1021/acs.jcim.5b00570>
  82. Polishchuk PG, Muratov EN, Artemenko AG et al (2009) Application of random forest approach to QSAR prediction of aquatic toxicity. *J Chem Inf Model* 49:2481–2488. <https://doi.org/10.1021/ci900203n>
  83. Marchese Robinson RL, Palczewska A, Palczewski J, Kidley N (2017) Comparison of the predictive performance and interpretability of random forest and linear models on benchmark data sets. *J Chem Inf Model* 57:1773–1792. <https://doi.org/10.1021/acs.jcim.6b00753>
  84. Cortes-Ciriano I, Murrell DS, van Westen GJP et al (2014) Prediction of the potency of mammalian cyclooxygenase inhibitors with ensemble proteochemometric modeling. *J Cheminf* 7:1. <https://doi.org/10.1186/s13321-014-0049-z>
  85. Cortes-Ciriano I, Firth NC, Bender A, Watson O (2018) Discovering highly potent molecules from an initial set of inactives using iterative screening. *J Chem Inf Model* 58:2000–2014. <https://doi.org/10.1021/acs.jcim.8b00376>
  86. Kalliokoski T, Kramer C, Vulpetti A, Gedeck P (2013) Comparability of mixed IC<sub>50</sub> data—a statistical analysis. *PLoS ONE* 8:e61007. <https://doi.org/10.1371/journal.pone.0061007>
  87. Koutsoukas A, Paricharak S, Galloway WRJD et al (2013) How diverse are diversity assessment methods? A comparative analysis and benchmarking of molecular descriptor space. *J Chem Inf Model* 54:230–242. <https://doi.org/10.1021/ci400469u>
  88. Todeschini R, Consonni V (2008) Handbook of molecular descriptors. (Wiley-VCH, 2000). <https://doi.org/10.1002/9783527613106>
  89. Kauvar LM, Higgins DL, Villar HO et al (1995) Predicting ligand binding to proteins by affinity fingerprinting. *Chem Biol* 2:107–118
  90. Fliri AF, Loging WT, Thadeio PF, Volkmann RA (2005) Biospectra analysis: model proteome characterizations for linking molecular structure and biological response. *J Med Chem* 48:6918–6925. <https://doi.org/10.1021/jm050494g>
  91. Fliri AF, Loging WT, Thadeio PF, Volkmann RA (2005) Biological spectra analysis: linking biological activity profiles to molecular structure. *Proc Natl Acad Sci USA* 102:261–266. <https://doi.org/10.1073/pnas.0407790101>
  92. Karimi M, Wu D, Wang Z, Shen Y (2019) DeepAffinity: interpretable deep learning of compound–protein affinity through unified recurrent and convolutional neural networks. *Bioinformatics*. <https://doi.org/10.1093/bioinformatics/btz111>
  93. Wallach I, Dzamba M, Heifets A (2015) AtomNet: a deep convolutional neural network for bioactivity prediction in structure-based drug discovery. <https://doi.org/10.1007/s10618-010-0175-9>. [arXiv:1510.02855](https://arxiv.org/abs/1510.02855). Accessed 8 July 2018
  94. Amidi A, Amidi S, Vlachakis D et al (2018) EnzyNet: enzyme classification using 3D convolutional neural networks on spatial representation. *PeerJ* 6:e4750. <https://doi.org/10.7717/peerj.4750>
  95. Derevyanko G, Grudin S, Bengio Y, Lamoureux G (2018) Deep convolutional networks for quality assessment of protein folds. *Bioinformatics* 34:4046–4053. <https://doi.org/10.1093/bioinformatics/bty494>

96. Torng W, Altman RB (2017) 3D deep convolutional neural networks for amino acid environment similarity analysis. *BMC Bioinform* 18:302. <https://doi.org/10.1186/s12859-017-1702-0>
97. Jiménez J, Škalič M, Martínez-Rosell G, De Fabritiis G (2018)  $K_{\text{DEEP}}$ : protein-ligand absolute binding affinity prediction via 3D-convolutional neural networks. *J Chem Inf Model* 58:287–296. <https://doi.org/10.1021/acs.jcim.7b00650>
98. Dubost F, Bortsova G, Adams H et al (2018) Hydranet: data augmentation for regression neural networks. [arXiv:1807.04798](https://arxiv.org/abs/1807.04798). Accessed 8 July 2018
99. Perez L, Wang J (2017) The effectiveness of data augmentation in image classification using deep learning. [arXiv:1712.04621](https://arxiv.org/abs/1712.04621). Accessed 8 July 2018
100. Taylor L, Nitschke G (2017) Improving deep learning using generic data augmentation. [arXiv:1708.06020](https://arxiv.org/abs/1708.06020). Accessed 8 July 2018
101. Bjerrum EJ (2017) SMILES enumeration as data augmentation for neural network modeling of molecules. [arXiv:1703.07076](https://arxiv.org/abs/1703.07076). Accessed 8 July 2018
102. Bansal M, Yang J, Karan C et al (2014) A community computational challenge to predict the activity of pairs of compounds. *Nat Biotechnol* 32:1213–1222. <https://doi.org/10.1038/nbt.3052>

### Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.