

MODELLING BIOLOGICAL ASSAYS WITH ADAPTIVE DEEP KERNEL LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

1 Due to the significant costs of data generation, many prediction tasks within drug
2 discovery are by nature few-shot regression (FSR) problems, including accurate
3 modelling of biological assays. Although a number of few-shot classification and
4 reinforcement learning methods exist for similar applications, we find relatively
5 few FSR methods meeting the performance standards required for such tasks under
6 real-world constraints. Inspired by deep kernel learning, we develop a novel FSR
7 algorithm that is better suited to these settings. Our algorithm consists of learning
8 a deep network in combination with a kernel function and a differentiable kernel
9 algorithm. As the choice of kernel is critical, our algorithm learns to find the
10 appropriate kernel for each task during inference. It thus performs more effectively
11 with complex task distributions, outperforming current state-of-the-art algorithms
12 on both toy and novel, real-world benchmarks that we introduce herein. By
13 introducing novel benchmarks derived from biological assays, we hope that the
14 community will progress towards the development of FSR algorithms suitable for
15 use in noisy and uncertain environments such as drug discovery.

16 1 INTRODUCTION

17 Following breakthroughs in domains including computer vision, autonomous driving, and natural
18 language processing, deep learning methods are now entering the domain of pharmaceutical R&D.
19 Recent successes include the deconvolution of biological targets from -omics data (Min et al.,
20 2017), generation of drug-like compounds via *de novo* molecular design (Xu et al., 2019), chemical
21 synthesis planning (Segler and Waller, 2017; Segler et al., 2017), and multi-modal image analysis for
22 quantification of cellular response (Min et al., 2017). A common characteristic of these applications,
23 however, is the availability of high quality, high quantity training data. Unfortunately, many critical
24 prediction tasks in the drug discovery pipeline fail to satisfy these requirements, in part due to
25 resource and cost constraints (Cherkasov et al., 2014).

26 We therefore focus this work on modelling biological assays (bio-assays) relevant in the early stages
27 of drug discovery, primarily binding and cellular readouts. Under the constraints of an active drug
28 discovery program, the data from these assays, consisting of libraries of molecules and their associated
29 real-valued activity scores, is often relatively small and noisy (refer to statistics in Section 5). In
30 many contexts, it can be challenging to build a training set of even a few dozen samples per individual
31 assay. Modelling an assay is thus best viewed as a few-shot regression (FSR) problem, with many
32 variables (including experimental conditions, readouts, concentrations, and instrument configurations)
33 accounting for the data distribution being generated. Practically, these variables make it infeasible to
34 compare data collected across different assays, thereby making it difficult to learn predictive models
35 from molecular structures. Furthermore, as bio-assay modelling is intended to be used for prioritizing
36 molecules for subsequent evaluation (e.g. Bayesian optimization) and efficiently exploring the overall
37 chemical space (e.g. active learning), accurate prediction and uncertainty estimation using few
38 datapoints is critical to successful application in drug discovery.

39 It is our view that robust FSR algorithms are needed to tackle this challenge. Specifically, we argue
40 that these algorithms should remain accurate in noisy environments, and also provide well-calibrated
41 uncertainty estimates to inform efficient exploration of chemical space during molecular optimization.
42 Fortunately, recent advances in few-shot learning have led to new algorithms that learn efficiently
43 and generalize adequately from small training data (Wang and Yao, 2019; Chen et al., 2019). Most

44 have adopted the meta-learning paradigm (Thrun and Pratt, 1998; Vilalta and Drissi, 2002), where
 45 some prior knowledge is learned across a large collection of tasks and then transferred to new tasks
 46 in which there are limited amounts of data. Such algorithms tend to differ in two aspects: the **nature**
 47 **of the meta-knowledge captured** and the **amount of adaptation performed at test-time** for new
 48 tasks or datasets. The meta-knowledge refers to the domain specific prior needed to solve each task
 49 most effectively. Due to the size of the total chemical space accessible when modelling bio-assays
 50 (Bohacek et al., 1996), there is a particular need for the meta-knowledge to be sufficiently rich so as
 51 to allow for extrapolation and uncertainty estimation in unseen regions of chemical space at test-time
 52 (i.e. for new tasks). Given that the same molecule can behave differently across different assays,
 53 greater test-time adaptation is also required and must be accounted for during modelling.

54 In previous work, metric learning methods (Koch et al., 2015; Vinyals et al., 2016; Snell et al.,
 55 2017; Garcia and Bruna, 2017; Bertinetto et al., 2018) accumulate meta-knowledge in high capacity
 56 covariance/distance functions and use simple base-learners such as k-nearest neighbor (Snell et al.,
 57 2017; Vinyals et al., 2016) or low capacity neural networks (Garcia and Bruna, 2017) to produce
 58 adequate models for new tasks. However, they do not adapt the covariance functions nor the base-
 59 learners at test-time. Initialization- and optimization-based methods (Finn et al., 2017; Kim et al.,
 60 2018; Ravi and Larochelle, 2016) that learn the initialization points and update rules for gradient
 61 descent-based algorithms, respectively, allow for improved adaptation on new tasks but remain time
 62 consuming and memory inefficient. We therefore argue that to ensure optimal performance when
 63 modelling bio-assays, it is crucial to combine the strengths of both types of methods while also
 64 allowing for the incorporation of domain-specific knowledge when making predictions. We achieve
 65 this by framing FSR as a deep kernel learning (DKL) task, deriving novel algorithms that we apply to
 66 modelling specific assays and readouts.

67 **Contributions:** Our contributions are three-fold. We first frame few-shot regression as a DKL
 68 problem and showcase its advantages relative to classical metric learning methods. We then derive
 69 the adaptive deep kernel learning (ADKL) framework by learning a conditional kernel function that
 70 is task dependant, allowing for more test-time adaptation than the DKL framework. Finally, we
 71 introduce two real-world datasets for modelling biological assays using FSR. With this contribution,
 72 we hope to encourage the development of subsequent few-shot regression methods suitable for
 73 real-world applications, as is the case for few-shot classification and reinforcement learning, each of
 74 which have received comparatively greater attention in recent years (Wang and Yao, 2019).

75 2 DEEP KERNEL LEARNING

76 In this section, we describe the DKL framework introduced for single tasks by Wilson et al. (2016).
 77 We then extend it to few-shot learning and discuss its advantages over metric learning algorithms.

78 **Single Task DKL:** Let $D_{trn}^t = \{(\mathbf{x}_i, y_i)\}_{i=1}^m \subset \mathcal{X} \times \mathbb{R}$, a training dataset available for learning
 79 the regression task t where \mathcal{X} is the input space and \mathbb{R} is the output space. A DKL algorithm aims
 80 to obtain a non-linear embedding of inputs in the embedding space \mathcal{H} , using a deep neural network
 81 $\phi_\theta : \mathcal{X} \rightarrow \mathcal{H}$ of parameters θ . It then finds the minimal norm regressor h_*^t in the reproducing kernel
 82 Hilbert space (RKHS) \mathcal{R} on \mathcal{H} , that minimize an objective function such as

$$h_*^t := \operatorname{argmin}_{h \in \mathcal{R}} \lambda \|h\|_{\mathcal{R}} + \ell(h, D_{trn}^t) \quad (1)$$

83 where ℓ is a non-negative loss function that measures the loss of a regressor h and λ weighs the
 84 importance of the norm minimization against the training loss. Following the representer theorem
 85 (Scholkopf and Smola, 2001; Steinwart and Christmann, 2008), h_*^t can be written as a finite linear
 86 combination of kernel evaluations on training inputs, i.e.:

$$h_*^t(\mathbf{x}) = \sum_{(\mathbf{x}_i, y_i) \in D_{trn}^t} \alpha_i^t k_\rho(\phi_\theta(\mathbf{x}), \phi_\theta(\mathbf{x}_i)), \quad (2)$$

87 where $\alpha^t = (\alpha_1^t, \dots, \alpha_m^t)$ are the learned combination weights and $k_\rho : \mathcal{H} \times \mathcal{H} \rightarrow \mathbb{R}_+$ is a
 88 chosen reproducing kernel of \mathcal{R} with hyperparameters ρ . Candidate kernels include the radial basis,
 89 polynomial, and linear kernels. α^t can be obtained by using a differentiable kernel method enabling
 90 the computation of the gradients of the loss w.r.t. the parameters θ . Such methods include Gaussian
 91 Process (GP), Kernel Ridge Regression (KRR), and Logistic Regression (LR).

92 As DKL inherits from deep learning and kernel methods, it follows that gradient descent algorithms are
 93 required to optimize the network parameters θ . The latter can be high dimensional and a substantial
 94 amount of training samples are required to train DKL models and avoid overfitting. However, once
 95 the latter condition is met, scalability of the kernel method can be limiting (running time in $O(m^3)$
 96 for m training samples) and approximations can be needed for scalability (see Williams and Seeger
 97 (2001); Wilson and Nickisch (2015)).

98 **Few-Shot DKL:** In few-shot learning, one has access to a meta-training collection $\mathcal{D}_{meta-trn} :=$
 99 $\left\{ (D_{trn}^{t_j}, D_{val}^{t_j}) \right\}_{j=1}^T$ of T tasks. Each task t_j has its own training (or support) set $D_{trn}^{t_j}$ and validation
 100 (or query) set $D_{val}^{t_j}$. A meta-testing collection $\mathcal{D}_{meta-tst}$ is also available to assess the generalization
 101 performance of the few-shot algorithm across unseen tasks. To obtain a Few-Shot DKL (FSDKL)
 102 method for FSR in such settings, one can share the parameters of ϕ_θ across all tasks, similar to metric
 103 learning algorithms. Hence, for a given task t_j , the inputs are first transformed by the function ϕ_θ
 104 and then a kernel method is used to obtain the regressor $h_*^{t_j}$, which will be evaluated on $D_{val}^{t_j}$. Here,
 105 KRR and GP are explored as they are the state-of-the-art algorithms for kernel-based regression. The
 106 latter is used to allow our models to provide accurate predictive uncertainty, which is useful when
 107 prioritizing molecules in the context of drug discovery.

KRR: Using the squared loss and the L2-norm to compute $\|h\|_{\mathcal{R}}$, KRR gives the optimal regressor
 for a task t and its validation loss $\mathcal{L}_{\theta, \rho, \lambda}^t$ as follows:

$$h_*^t(\mathbf{x}) = \alpha K_{\mathbf{x}, trn}, \quad \text{with} \quad \alpha = (K_{trn, trn} + \lambda I)^{-1} \mathbf{y}_{trn} \quad (3)$$

$$\mathcal{L}_{\theta, \rho, \lambda}^t = \mathbf{E}_{\mathbf{x}, y \sim D_{val}^t} (\alpha K_{\mathbf{x}, trn} - y)^2, \quad (4)$$

108 where $\mathbf{y}_{trn} = (y_1, \dots, y_{|D_{trn}^t|})^T$, $K_{trn, trn}$ is the matrix of kernel evaluations where entry i, l is
 109 $k_\rho(\phi_\theta(\mathbf{x}_i), \phi_\theta(\mathbf{x}_l))$ for pairs of examples in D_{trn}^t . An equivalent definition applies to $K_{\mathbf{x}, trn}$.

GP: When using the negative log likelihood loss function, the GP algorithm gives a probabilistic
 regressor for which the predictive mean, covariance, and loss for a task t are:

$$\mathcal{L}_{\theta, \rho, \lambda}^t = -\ln \mathcal{N}(\mathbf{y}_{val}; \mathbb{E}[h_*^t], \text{cov}(h_*^t)), \quad (5)$$

$$\mathbb{E}[h_*^t] = K_{val, trn} (K_{trn, trn} + \lambda I)^{-1} \mathbf{y}_{trn}, \quad (6)$$

$$\text{cov}(h_*^t) = K_{val, val} - K_{val, trn} (K_{trn, trn} + \lambda I)^{-1} K_{trn, val} \quad (7)$$

110 Finally, the parameters θ of the neural network, along with λ and the kernel hyperparameters ρ , are
 111 optimized using the expected loss on all tasks:

$$\underset{\theta, \rho, \lambda}{\text{argmin}} \mathbf{E}_{t \sim \mathcal{D}_{meta-trn}} \mathcal{L}_{\theta, \rho, \lambda}^t. \quad (8)$$

112 To summarize, FSDKL finds a representation common to all tasks such that the kernel method (in our
 113 case, GP and KRR) will generalize well from a small amount of samples. In doing so, this alleviates
 114 two of the main limitations of single task DKL: i) the scalability of the kernel method is no longer an
 115 issue since we are in the few-shot learning regime¹, and ii) the parameters θ (and ρ, λ) are learned
 116 across a potentially large amount of tasks and samples, providing the opportunity to learn a rich
 117 representation without overfitting.

118 Despite shared characteristics with the metric learning framework, the FSDKL framework is more
 119 powerful and flexible. It provides better task-specific adaptation due to the inference of the appropriate
 120 model using the kernel methods compared to shared model parameters in metric learning. After meta-
 121 training, any task-specific model also inherits the generalization guarantees of kernel-based models,
 122 and consequently increasing the number of shots for new tasks can only improve generalization
 123 performance. The incorporation of prior knowledge through user-specific kernel functions is also
 124 a major advantage of DKL over metric learning (e.g. use periodic kernels for periodic function
 125 regression tasks).

¹Even with several hundred samples, the computational cost of embedding each example is usually higher than inverting the Gram matrix.

126 3 ADAPTIVE DEEP KERNEL LEARNING

127 In this section, we present a new algorithm, deemed adaptive deep kernel learning (ADKL). Funda-
 128 mentally, it differs from FSDKL by having more flexibility in its kernel definition and by learning to
 129 produce task-specific kernel functions during the meta-training instead of using one defined by the
 130 user. It does so by learning a task representation using a task encoding network ψ_η and leveraging it
 131 to build task-specific kernels using a multi-modal neural network c_ρ . More explicitly, given a task t ,
 132 ADKL first computes a task embedding $\mathbf{z}_t = \psi_\eta(D_{trn}^t)$ using its support set D_{trn}^t and then it infers
 133 the adapted kernel with c_ρ . We describe in more detail both the task encoding network ψ_η and the
 134 network c_ρ responsible for computing the task-specific kernel below.




Figure 1: ADKL-KRR. The blue and orange colors show the procedure for a task during internal train and test, respectively. During training, ADKL first computes a task embedding $\mathbf{z}_t = \psi_\eta(D_{trn}^t)$ that is used with a pseudo-representations U by the network c_ρ to produce the task-specific kernel function. The empirical kernel map of this kernel gives the function $C_t(\cdot)$ that is evaluated for every training point to produce $K_{trn,trn}$. The latter and the train targets are used by KRR (or GP) to produce the model h_t^* . At evaluation, $C_t(\cdot)$ is evaluated again for every test point to obtain $K_{val,trn}$, which is used to compute the predictions. The loss is then computed and used to update all parameters of ADKL.

135 3.1 TASK ENCODING

136 The challenge of the network ψ_η is to capture complex dependencies in the training set D_{trn}^t to
 137 provide a useful task encoding \mathbf{z}_t . Furthermore, the task encoder should be invariant to permutations
 138 of the training set and be able to encode a variable amount of samples. After exploring a variety
 139 of architectures, we found that those that are more complex, such as Transformers (Vaswani et al.,
 140 2017), tend to underperform. This is possibly due to overfitting or the sensitivity of training such
 141 architectures.

142 Consequently, inspired by DeepSets (Zaheer et al., 2017), we propose a simple order invariant
 143 network that captures the first and second order statistics of regression datasets. Given a dataset, this
 144 network first processes each of its samples individually as follows: a) extract input features using
 145 ϕ_θ (see section 2), b) concatenate the input features with the target and embed the obtained vector
 146 using a simple fully connected neural network \mathbf{r}_η of parameters η ². It then computes the first and
 147 the second order statistics of the obtained vectors for all samples of the dataset and concatenates them
 148 to produce the representation. More formally,

$$\psi_\eta(D_{trn}^t) := [\mu^t, \sigma^t], \text{ with } \mu^t = \mathbf{E}_{(\mathbf{x}, y) \in D_{trn}^t} \mathbf{r}_\eta([\phi_\theta(\mathbf{x}_i), y_i]), \quad \sigma^t = \mathbf{Var}_{(\mathbf{x}, y) \in D_{trn}^t} \mathbf{r}_\eta([\phi_\theta(\mathbf{x}_i), y_i])$$

149 where $[\cdot, \cdot]$ is the concatenation operator. As μ^t and σ^t are invariant to permutations in D_{trn}^t , it
 150 follows that ψ_η is also permutation invariant. Overall, ψ_η is simply the concatenation of the first and
 151 second moments of the sample representations, which were nonlinear transformations of the original
 152 inputs and targets.

To help the training of the parameters η , we add a regularization term that maximizes the mutual information between D_{trn}^t and D_{val}^t . This encourage the network to produce similar task encodings

²These are the only parameters involved in the computation of the task encoding, which is why we also use the notation ψ_η .

when presented with different data partitions for a given task. Concretely, we maximize the lower bound on the mutual information between the task representations given by the support and the query sets instead of the true mutual information (Belghazi et al., 2018). Using a batch of b tasks and the cosine similarity c as the similarity measure between two task encodings, this lower bound \tilde{I}_η is defined by Eq. (9) and is the regularizer that we used to have a better task encoder.

$$\tilde{I}_\eta \stackrel{\text{def}}{=} \frac{1}{b} \sum_{j=1}^b c(\psi_\eta(D_{trn}^{t_j}), \psi_\eta(D_{val}^{t_j})) - \ln \frac{1}{b(b-1)} \sum_{j=1}^b \sum_{i \neq j} c^c(\psi_\eta(D_{trn}^{t_j}), \psi_\eta(D_{val}^{t_i})) \quad (9)$$

153 3.2 TASK-SPECIFIC KERNEL

154 Here, we describe how the task-specific kernels are inferred using the task representations described
155 previously. In fact, they are all obtained using a multi-modal neural network c_ρ of parameter ρ .
156 Given any pair of input representations $(\phi(\mathbf{x})$ and $\phi(\mathbf{x}'))$ and a task encoding \mathbf{z}_t , this network simply
157 computes the input pair similarity under the condition given by the task encoding as follows:

$$c_\rho(\phi(\mathbf{x}), \phi(\mathbf{x}'), \mathbf{z}_t) := MLP_\rho([\phi(\mathbf{x}) - \phi(\mathbf{x}')]^2, \mathbf{z}_t], \quad (10)$$

158 where $(\phi(\mathbf{x}) - \phi(\mathbf{x}'))^2$ is the element-wise L2 distance between the input representations, $[\cdot, \cdot]$ is
159 the concatenation operator and MLP is a fully connected neural network of parameters ρ which has a
160 single neuron at its last layer. It bears mentioning that c_ρ is symmetric and stationary with regard
161 to $\phi(\mathbf{x})$ and $\phi(\mathbf{x}')$ as their element-wise L2 distances vector is received as part of the input of the
162 fully connected network. Further, by simply concatenating the task representation \mathbf{z}_t to this distance
163 vector at the input, c_ρ provides a powerful approach to produce task-specific kernels. However, these
164 kernels are not positive semi-definite (PSD) and cannot be directly used for KRR and GP. Therefore,
165 using the empirical kernel mapping technique (Schölkopf et al., 1999) we computed the task-specific
166 PSD kernel $k_{\rho,t}$ associated with a given task representation \mathbf{z}_t obtained from D_{trn}^t . This kernel can
167 be written as the empirical kernel map of $c_\rho(\cdot, \cdot, \mathbf{z}_t)$ with regard to D_{trn}^t i.e.:

$$k_{\rho,t}(\mathbf{x}, \mathbf{x}') = C_t(\mathbf{x}) \cdot C_t(\mathbf{x}'), \quad \text{with} \quad (11)$$

$$C_t(\mathbf{x}) = (c_\rho(\mathbf{x}, \mathbf{x}_1, \mathbf{z}_t), \dots, c_\rho(\mathbf{x}, \mathbf{x}_m, \mathbf{z}_t)), \quad \text{and} \quad (\mathbf{x}_i, \cdot) \in D_{trn}^t \forall i = 1, \dots, m$$

168 Using the empirical kernel map of c_ρ to compute $k_{\rho,t}$ offers the opportunity to introduce *pseudo-input*
169 *representations* (or *pseudo-representations*) that could improve the kernel evaluations, specially in
170 low data settings. More precisely, instead of computing the empirical kernel map with regard to
171 D_{trn}^t alone, we use $(D_{trn}^t \cup U)$ where U is the set of pseudo-representations. The function C_t , from
172 Eq. (11), becomes:

$$C_t(\mathbf{x}) = (c_\rho(\mathbf{x}, \mathbf{x}_1, \mathbf{z}_t), \dots, c_\rho(\mathbf{x}, \mathbf{x}_m, \mathbf{z}_t), c_\rho(\mathbf{x}, \mathbf{u}_1, \mathbf{z}_t), \dots, c_\rho(\mathbf{x}, \mathbf{u}_l, \mathbf{z}_t)), \quad (12)$$

$$\text{with} \quad \mathbf{u}_l \in U \forall l = 1, \dots, |U| \quad \text{and} \quad (\mathbf{x}_i, \cdot) \in D_{trn}^t \forall i = 1, \dots, m$$

173 The number of pseudo-representations is a hyperparameter of ADKL (in our experiments we choose
174 $|U| \in [0, 50]$) and all pseudo-representations $\mathbf{u}_l \in \mathcal{H}$ are learnable parameters that are shared by
175 all tasks and learned during meta-training. To prevent their collapse into a single point during the
176 training and to ensure that they are well distributed in the feature space \mathcal{H} , we add a regularization
177 term (\tilde{D}_u) to the training loss function. To introduce this regularization term, let's consider p and
178 q to be the distributions that generate the true input representations and the pseudo-representations,
179 respectively. We make the assumption that p and q are both multivariate Gaussian distributions with
180 diagonal covariance matrices and have respective parameters $(\mu_\phi, \sigma_\phi^2)$ and (μ_u, σ_u^2) . The parameters
181 of p are estimated using the running means and variances of all input representations computed over
182 batches of tasks. Those of q are estimated using U . As, p and q must be close, the training of the
183 pseudo-representations is regularized by minimizing the KL distance \tilde{D}_u between p and q , i.e.:

$$\tilde{D}_u = KL(\mathcal{N}(\mu_u, \sigma_u^2) \parallel \mathcal{N}(\mu_\phi, \sigma_\phi^2)) \quad (13)$$

Putting it all together, the ADKL training objective is the following:

$$\underset{\theta, \eta, \rho, \mathbf{u}, \lambda}{\text{argmin}} \quad \mathbf{E}_{t_j \sim B} \mathcal{L}_{\theta, \eta, \rho, \mathbf{u}, \lambda}^{t_j} - \gamma_{task} \tilde{I}_\eta + \gamma_{pseudo} \tilde{D}_u, \quad (14)$$

184 with $\gamma_{task} \geq 0$ as a tradeoff hyperparameter for the regularization of the task-encoder and $\gamma_{pseudo} \geq 0$
185 as a tradeoff hyperparameter for the regularization of the pseudo-inputs.

186 4 RELATED WORK

187 Across the spectrum of learning approaches, DKL methods lie between neural networks and kernel
 188 methods. While neural networks can learn from a very large amount of data without much prior
 189 knowledge, kernel methods learn from fewer data when given an appropriate covariance function
 190 that accounts for prior knowledge of the relevant task. In the first DKL attempt, Wilson et al. (2016)
 191 combined GP with CNN to learn a covariance function adapted to a task from large amounts of data,
 192 though the large time and space complexity of kernel methods forced the approximation of the exact
 193 kernel using KISS-GP (Wilson and Nickisch, 2015). Dasgupta et al. (2018) have demonstrated that
 194 such approximation is not necessary using finite rank kernels. Here, we show that learning from a
 195 collection of tasks (FSR mode) does not require any approximation when the covariance function is
 196 shared across tasks. This is an important distinction between our study and other existing studies in
 197 DKL, which learn their kernel for single task applications instead of multiple task collections.

198 On the spectrum between NNs and kernel methods we must also mention metric learning. Metric
 199 learning algorithms learn an input covariance function shared across tasks but rely only on the
 200 expressive power of DNNs. First, stochastic kernels are built out of shared feature extractors and
 201 simple pairwise metrics (e.g. cosine similarity (Vinyals et al., 2016), Euclidean distance (Snell et al.,
 202 2017)), or parametric functions (e.g. relation modules (Sung et al., 2018), graph neural networks
 203 (Garcia and Bruna, 2017; Kim et al., 2019a)). Then, within tasks, the predictions are distance-
 204 weighted combinations of the training labels with the stochastic kernel evaluations—no adaptation is
 205 done.

206 In connection with the test-time adaptation capabilities of our method, methods that combine metric
 207 learning with initialization-based models are great competitors. In fact, Proto-MAML (Triantafillou
 208 et al., 2019), which captures the best of Prototypical Networks (Snell et al., 2017) and MAML
 209 (Finn et al., 2017), allows within-task adaptation using MAML on top of a shared feature extractor.
 210 Similarly, Kim et al. (2018) have proposed a Bayesian version of MAML where a feature extractor is
 211 shared across tasks, while multiple MAML particles are used for the task-level adaptation. Bertinetto
 212 et al. (2018) have also tackled the lack of adaptation for new tasks by using Ridge Regression and
 213 Logistic Regression to find the appropriate weighting of the training samples for classification tasks.
 214 This study can be considered as an instance of the FSDKL framework, though its contribution was
 215 limited to showing that simple differentiable learning algorithms can increase adaptation in the metric
 216 learning framework. Our work goes beyond by formalizing few-shot DKL and proposing ADKL: a
 217 data-driven manner to compute the correct kernel for a task.

218 Since ADKL-GP learns task-specific stochastic processes, it is related to neural processes (Garnelo
 219 et al., 2018a) and the ML-PIP framework (Gordon et al., 2018). Both propose a scalable alternative
 220 to learning regression functions by performing inference on stochastic processes. In these families
 221 of methods, both Conditional Neural Processes (CNP) (Garnelo et al., 2018b) and Attentive Neural
 222 Processes (ANP) (Kim et al., 2019b) learn conditional stochastic processes parameterized by task-
 223 specific conditions derived from the support sets, but CNP is the most related to ADKL-GP. CNP is
 224 an instance of ML-PIP when the task encoder gives a point estimate of the task parameters instead
 225 of a distribution. Finally, the main differences between ANP and CNP are the architecture of the
 226 task-encoder and the lack of mathematical guarantees associated with stochastic processes in CNP
 227 (as it does not impose any consistency with respect to a prior process). By comparison, ADKL-GP
 228 also learns conditional stochastic processes but has mathematical guarantees thanks to GP and PSD
 229 kernels.

230 5 DATASETS

231 Existing FSR methods have been mostly tested on 1D function regression and pixel-wise image
 232 completion tasks with MNIST and CelebA (Kim et al., 2018; Garnelo et al., 2018b;a). On one hand,
 233 the 1D regression tasks are all relatively simple, almost noise-less, and homogeneous. On the other
 234 hand, methods have been successful for image completion tasks only outside the few-shot regime (i.e.
 235 when the number of samples is greater than 500) (Garnelo et al., 2018b;a). For these reasons, we
 236 introduce two task collections from a real-world context. Deemed **Binding** and **Antibacterial**, these
 237 task collections contain data from bio-assays that are representative of real-world FSR tasks in drug

238 discovery. The pre-processed versions of these collections and detailed statistics are available [here](#)
239 ([anonymized link](#)).

240 **Binding:** All tasks in this collection aim to predict the binding affinity of small molecules to a target
241 protein. The characteristics of the proteins thus define different data distributions over the chemical
242 space. The inputs and the targets for each task are molecules that have been tested in a binding
243 assay and the measured binding affinity of the molecule against a given protein. The task collection
244 was extracted from the public database [BindingDB](#) and altered by removing bio-assays with targets
245 correlated above 0.8 or those with less than 10 experimental measurements, leaving us with 5, 717
246 tasks.

247 **Antibacterial:** Within this collection, the task is to predict the antimicrobial activity of small
248 molecules against various bacteria. They are characterized by a bacterial strain whose resistance to
249 drug-like molecules was being evaluated. The task collection was extracted from the public database
250 [PubChem](#). After also removing bio-assays with correlations above 0.8 and those with less than 10
251 samples, we obtain 3, 255 tasks.

252 Their meta-test partitions each contain 1000 tasks, with the remaining used in the meta-train and
253 meta-validation. The molecules (represented as [SMILES](#)) are converted into vectors using routines
254 available in the RDKit software (more precisely into [ECFP6](#) binary fingerprint vectors of 4,096
255 dimensions). These inputs were also processed in all methods using the same feature extractor
256 architecture, which is a fully-connected network of $256 \times 256 \times 256$. Due to the high noise-to-signal
257 ratio, the targets are first \log_2 -scaled and then scaled linearly between 0 and 1 to avoid scaling issues
258 during training.

259 [Fig. 2](#) highlights three aspects of the collections that make them complementary to existing bench-
260 marks, but better suited for evaluating the readiness of FSR methods for real-world applications
261 relative to toy collections. First, the distributions of number of samples per task show that they
262 naturally contain few samples, which we believe reflects the costs of acquiring labelled data in a
263 drug discovery setting. In comparison, the number of samples available per task is relatively large in
264 previous benchmarks, with the few-shot regime being achieved artificially through sampling. Second,
265 as illustrated by their noise-to-signal ratio, real-world tasks are inherently noisy, increasing the
266 difficulties associated with few-shot learning. Finally, the input diversity within each task is reduced
267 relative to the total among tasks. Despite this diversity difference, good models should perform
268 relatively well outside the input region they have seen in the support set. This situation challenges the
269 methods to learn strong priors about the input space and to be able to generalize after seeing only a
270 small fraction of it. These collections invite researchers to explore meta-learning with increasingly
271 heterogeneous datasets and in noisy environments, as well as generalisation and extrapolation in
272 large input spaces (such as the drug-like chemical space, which is estimated to be approximately 10^{33}
273 molecules (Polishchuk et al., 2013)).

274 To test our method in a noise-less environment, we also use the **Sinusoids** collection introduced by
275 Kim et al. (2018). This challenging few-shot regression benchmark consists of 5,000 tasks defined
276 by functions of the form: $y = A \sin(wx + b) + \epsilon$ with $A \in [0.1, 5.0]$, $b \in [0.0, 2\pi]$, and
277 $w \in [0.5, 2.0]$. Sampling inputs $x \in [-5.0, 5.0]$ and observational noise $\epsilon \in N(0, (0.01A)^2)$ and
278 computing y gives the samples for each task. Here, the meta-train, meta-validation, and meta-test
279 contain 56.25%, 18.75% and 25% of all the tasks, respectively, and all methods use the same feature
280 extractor architecture, which is a fully-connected network of $40 \times 40 \times 40$.

281 6 EXPERIMENTS

282 6.1 BENCHMARKING ANALYSIS

283 For all benchmarks, the performances of ADKL is compared against other meta-learning algorithms:
284 R2-D2 (an instance of FSDKL Bertinetto et al. (2018)), CNP (Garnelo et al., 2018b), MAML (Finn
285 et al., 2017), BMAML (Kim et al., 2018), and Learned Basis (Yi Loo, 2019) (all implementations are
286 available [here](#) ([anonymized link](#))). These algorithms have all proven to have efficient and effective
287 test-time adaptation routines and therefore constitute strong baselines for benchmarking. However,
288 for bioassay modelling benchmarks, we have also added two methods considered to be state-of-the-art
289 in chemoinformatics to assess performance relative to all meta-learning approaches. These methods




Figure 2: Statistics on bio-assay modelling tasks. Left: Number of samples per task. Middle: Noise-to-signal ratio. Right: Within-task versus overall molecular diversity.

k	5	10	20
ADKL-GP			
ADKL-KRR	0.0380 ± 0.0020	0.0348 ± 0.0009	0.0322 ± 0.0020
BMAML	0.0813 ± 0.0571	0.0486 ± 0.0071	0.0487 ± 0.0012
CNP	0.0416 ± 0.0019	0.0393 ± 0.0030	0.0397 ± 0.0027
ECFP4+KRR	0.0376 ± 0.0012	0.0352 ± 0.0014	0.0317 ± 0.0016
ECFP4+RF	0.0373 ± 0.0012	0.0339 ± 0.0013	0.0311 ± 0.0012
LearnedBasis	0.0761 ± 0.0040	0.0754 ± 0.0042	0.0616 ± 0.0215
R2D2	0.0492 ± 0.0015	0.0460 ± 0.0110	0.0342 ± 0.0012

Table 1: Average MSE on Binding

k	5	10	20
ADKL-GP	0.1017 ± 0.0013	0.0895 ± 0.0015	0.0860 ± 0.0016
ADKL-KRR	0.1000 ± 0.0012	0.0893 ± 0.0015	0.0862 ± 0.0009
BMAML	0.1059 ± 0.0021	0.1020 ± 0.0029	0.4616 ± 0.4210
CNP	0.1063 ± 0.0023	0.1239 ± 0.0219	0.1382 ± 0.0049
ECFP4+KRR	0.1166 ± 0.0020	0.1003 ± 0.0009	0.0956 ± 0.0009
ECFP4+RF	0.1129 ± 0.0002	0.1016 ± 0.0008	0.0970 ± 0.0003
LearnedBasis	0.1274 ± 0.0037	0.1308 ± 0.0032	0.1329 ± 0.0043
R2D2	0.1104 ± 0.0023	0.0962 ± 0.0021	0.0921 ± 0.0010

Table 2: Average MSE on Antibacterial

290 are the Random Forest algorithm with ECFP4 (Extended Connectivity FingerPrints of diameter 4) as
 291 molecular input representation, and ECFP4 with KRR and tanimoto similarity as a kernel function
 292 (Olier et al., 2018). During meta-test, each task is partitioned into query and support sets, then the
 293 support set is used to generate a model which is evaluated on the query set to compute the MSE. This
 294 process is repeated 30 times per task and the average MSE over the repetitions per task and over all
 295 tasks is reported in Tables 1 to 3.

296 For the Sinusoids collection, Table 3 shows that DKL-based methods significantly outperform all
 297 other methods despite their test-time adaptation capabilities. These results alone demonstrate the
 298 effectiveness of DKL-based methods in FSR relative to the current state-of-the-art. Furthermore, of
 299 all DKL-based methods, ADKL-KRR shows consistently stronger performance than others. This
 300 demonstrates that using ADKL increases test-time performance relative to FS-DKL (as R2-D2 and
 301 ADKL-KRR only differ by the kernel definition). It also indicates that attempting to capture the
 302 model uncertainty using GP in ADKL (instead of KRR) comes with a significant cost, especially in
 303 lower data regimes. This may be due to the inability of GP to differentiate between the observational
 304 noise and the model uncertainty as the number of samples get smaller. Also, notice that all task
 305 encoding based methods significantly outperform the others. This shows that adequately capturing
 306 the task representation is crucial for this task collection, and ADKL-KRR appears to be best equipped
 307 to do so.

308 Tables 1 and 2 show the performances of all methods on real-world datasets. As complements,
 309 Tables 4 and 5 show the p -value that assesses the statistical significance of the difference between
 310 each model and ADKL-GP and ADKL-KRR. These p -values result from Wilcoxon ranked tests
 311 comparing the MSE per task of each algorithm to ADKLs. Combined together, these tables shows
 312 that ADKL methods significantly outperforms all other meta-learning methods (p -values < 0.05).
 313 They also outperformed the state-of the art in chemoinformatic for Antibacterial, but do not on
 314 BINDING where those methods are significantly better than all meta-learning algorithms. Even if,
 315 ADKL is a first step in the right direction, these results show that there remains much room to develop
 316 meta-learning algorithms which are undoubtedly superior to methods in computational chemistry. It
 317 is also worth noticing that ADKL methods are significantly better than R2-D2 for these collections
 318 also confirming that using task specific kernels are useful and improve generalization.

m	5	10	20
model			
BMAML	2.042	1.371	0.844
CNP	1.616	0.392	0.117
Learned Basis	3.587	0.800	0.127
MAML	2.896	1.634	0.901
ADKL-GP	1.178	0.084	0.007
ADKL-KRR	0.867	0.061	0.005
FSDKL (R2D2)	1.002	0.073	0.009

Table 3: Average MSE on Sinusoidals

	ADKL-GP	ADKL-KRR
ADKL-GP		1.21e-01
CNP	0.00e+00	0.00e+00
ADKL-KRR	1.21e-01	
ECFP4+KRR	2.48e-78	3.69e-62
LearnedBasis	0.00e+00	0.00e+00
BMAML	0.00e+00	0.00e+00
FSDKL (R2D2)	1.18e-41	5.50e-15
ECFP4+RF	9.70e-81	2.20e-168

Table 4: Wilcoxon p-values – BindingDB

	ADKL-KRR
ADKL-GP	
CNP	1.77e-114
ADKL-KRR	
ECFP4+KRR	2.70e-03
LearnedBasis	0.00e+00
BMAML	0.00e+00
FSDKL (R2D2)	1.76e-49
ECFP4+RF	5.94e-01

Table 5: Wilcoxon p-values – BindingDB

319 6.2 ACTIVE LEARNING

320 In this section, we report the results of active learning experiments. Our intent is to measure the
 321 effectiveness of the uncertainty captured by the predictive distribution of ADKL-GP for active
 322 learning, as it is critical to our drug discovery use-cases. CNP, in comparison, serves to measure
 323 which of CNP and GP better captures the data uncertainty for improving FSR under active sample
 324 selection. For this purpose, we meta-train both algorithms using support and query sets of size
 325 $m = 5$. During meta-test time, five samples are randomly selected to constitute the support set D_{trn}
 326 and build the initial hypothesis for each task. Then, from a pool U of unlabeled data, we choose the
 327 input \mathbf{x}^* of maximum predictive entropy, i.e. $\mathbf{x}^* = \operatorname{argmax}_{\mathbf{x} \in U} \mathbb{E} [\log p(y|\mathbf{x}, D_{trn})]$. The latter is
 328 removed from U and added to D_{trn} with its predicted label. The within-task adaptation is performed
 329 on the new support set to obtain a new hypothesis which is evaluated on the query set D_{val} of the
 330 task. This process is repeated until we reach the allotted budget of 20 queries.

331 Fig. 3 illustrates, for all collections, the MSE after each sample acquisition iteration and under both
 332 random and active learning acquisition strategies. Under the active learning strategy, ADKL-GP
 333 consistently outperforms CNP. In particular, we observe that very few samples are queried by ADKL-GP
 334 to capture the data distribution whereas CNP performance remains far from optimal even when
 335 allowed the maximum number of queries. Further, since using the maximum predictive entropy
 336 strategy is better than querying samples at random for ADKL-GP (solid vs. dashed line), these results
 337 suggest that the predictive uncertainty obtained with GP is informative and more accurate than that of
 338 CNP. Moreover, when the number of queries is greater than 10, we observe a performance degradation
 339 for CNP while ADKL-GP remains consistent. This observation highlights the generalization capacity
 340 of DKL methods, even outside the few-shot regime where they have been trained — this same
 341 property does not hold true for CNP. We attribute this property of DKL methods to their use of kernel
 342 methods. In fact, their role in adaptation and generalization increases as we move away from the
 343 few-shot training regime.




Figure 3: Average MSE performance on the meta-test during active learning. The width of the shaded regions denotes the uncertainty over five runs for the sinusoidal collection. No uncertainty is shown for the real-world tasks as they were too time consuming.

344 6.3 ABLATION EXPERIMENTS

345 In our final set of experiments, we more closely evaluate the impact of the task encoder and the pseudo-
 346 inputs on the generalization during meta-testing. We do so by training and evaluating ADKL on

347 Sinusoids with different hyperparameter combinations. Figs. 4a to 4d show the relative improvements
 348 (negative values) or setbacks (positive values) in the meta-test MSE compared to different baselines
 349 (but the joint impact of γ_{task} and γ_{pseudo} is only discussed in Appendix A.3).

350 First, Fig. 4a compares $\gamma_{task} \in \{0.01, 0.1\}$ relative to $\gamma_{task} = 0$ and consequently demonstrates that
 351 regularizing the task encoder by maximizing the mutual information between the support set and
 352 the query set significantly improves the generalization performance. This conclusion holds for all
 353 support set sizes tested, as shown in Appendix A.1. Combined with the results from Section 6.1, this
 354 figure shows the importance of good task encoders for generalization in few-shot learning and how
 355 using the regularization term that we introduced is a step forward in that direction.

356 Then, Fig. 4c measures the relative differences between $\gamma_{pseudo} \in \{0.01, 0.1\}$ and $\gamma_{pseudo} = 0$ for
 357 different values of hyperparameter combinations. It shows that improving the kernel map evaluations
 358 using *pseudo-input representations* can significantly help with the generalization performance of
 359 ADKL. This conclusion also holds for all values tested for $|D_{trn}^t|$ (see Appendix A.2). However, the
 360 improvements were more consistent for smaller support sets, which is not surprising as improving
 361 the kernel map estimations in these cases is more critical.

362 Finally, Figs. 4b and 4d illustrate for ADKL-GP and ADKL-KRR, and different sizes of support sets,
 363 how the number of pseudo-representations (i.e $|U|$) affects performance. The values for each cell
 364 are relative performance using $|U| \in \{20, 50\}$ versus $|U| = 0$ and have been averaged over different
 365 hyperparameters and γ_{pseudo} . In general, we can confirm that increasing the number of pseudo-
 366 representations increases the estimates of the kernel maps and improves generalization. However, the
 367 improvements are more prominent with KRR in comparison to GP, which may be due to the fact that
 368 GP attributes a part of the modelling noise to the kernel evaluations, leading to more constraints on
 369 the optimization of the pseudo-representation parameters.

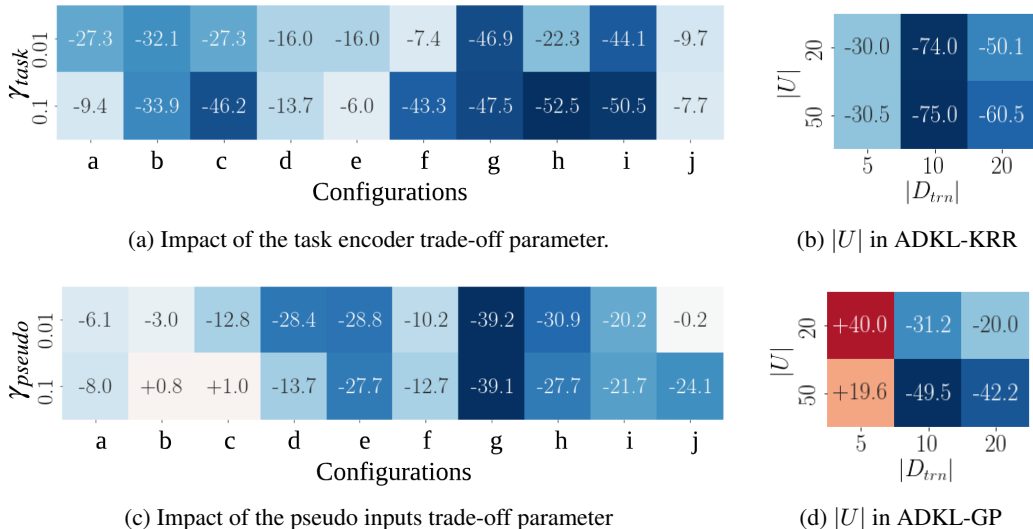


Figure 4: Relative decrease/increase in the meta-test MSE compared to different baselines. In (a) and (c) the baselines are $\gamma_{task} = 0$ and $\gamma_{pseudo} = 0$, respectively. In (b) and (d) the baselines are $|U| = 0$

370 **7 CONCLUSION**

371 We investigate bio-assays modelling using FSR methods. Our proposed method, ADKL, stores
 372 meta-knowledge in kernel functions and adapts to new tasks using KRR or GP. Our experiments
 373 provide evidence that the additional adaptation capacity at test-time provided by ADKL increases
 374 generalization significantly. Also, in a Bayesian setup, ADKL provides better predictive uncertainty,
 375 increasing their utility in bioassay modelling. However, there is still room to improve ADKL and
 376 most meta-learning methods to be better than traditional chemoinformatic methods. We hope, by
 377 making our bio-assay task collections publicly available, that the community will leverage them to
 378 propose new competitive FSR methods.

379 REFERENCES

- 380 Seonwoo Min, Byunghan Lee, and Sungroh Yoon. Deep learning in bioinformatics. *Briefings in*
381 *bioinformatics*, 18(5):851–869, 2017.
- 382 Youjun Xu, Kangjie Lin, Shiwei Wang, Lei Wang, Chenjing Cai, Chen Song, Luhua Lai, and Jianfeng
383 Pei. Deep learning for molecular generation. *Future medicinal chemistry*, 11(6):567–597, 2019.
- 384 Marwin HS Segler and Mark P Waller. Neural-symbolic machine learning for retrosynthesis and
385 reaction prediction. *Chemistry—A European Journal*, 23(25):5966–5971, 2017.
- 386 Marwin HS Segler, Mike Preuss, and Mark P Waller. Learning to plan chemical syntheses. *arXiv*
387 *preprint arXiv:1708.04202*, 2017.
- 388 Artem Cherkasov, Eugene N Muratov, Denis Fourches, Alexandre Varnek, Igor I Baskin, Mark
389 Cronin, John Dearden, Paola Gramatica, Yvonne C Martin, Roberto Todeschini, et al. Qsar
390 modeling: where have you been? where are you going to? *Journal of medicinal chemistry*, 57(12):
391 4977–5010, 2014.
- 392 Yaqing Wang and Quanming Yao. Few-shot learning: A survey. *CoRR*, abs/1904.05046, 2019. URL
393 <http://arxiv.org/abs/1904.05046>.
- 394 Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang. A closer look
395 at few-shot classification. *arXiv preprint arXiv:1904.04232*, 2019.
- 396 Sebastian Thrun and Lorien Pratt. Learning to learn: Introduction and overview. In *Learning to learn*,
397 pages 3–17. Springer, 1998.
- 398 Ricardo Vilalta and Youssef Drissi. A perspective view and survey of meta-learning. *Artificial*
399 *intelligence review*, 18(2):77–95, 2002.
- 400 Regine S Bohacek, Colin McMartin, and Wayne C Guida. The art and practice of structure-based
401 drug design: a molecular modeling perspective. *Medicinal research reviews*, 16(1):3–50, 1996.
- 402 Gregory Koch, Richard Zemel, and Ruslan Salakhutdinov. Siamese neural networks for one-shot
403 image recognition. In *ICML deep learning workshop*, volume 2, 2015.
- 404 Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one
405 shot learning. In *Advances in neural information processing systems*, pages 3630–3638, 2016.
- 406 Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In
407 *Advances in Neural Information Processing Systems*, pages 4077–4087, 2017.
- 408 Victor Garcia and Joan Bruna. Few-shot learning with graph neural networks. *arXiv preprint*
409 *arXiv:1711.04043*, 2017.
- 410 Luca Bertinetto, Joao F Henriques, Philip HS Torr, and Andrea Vedaldi. Meta-learning with differen-
411 tiable closed-form solvers. *arXiv preprint arXiv:1805.08136*, 2018.
- 412 Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of
413 deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume*
414 *70*, pages 1126–1135. JMLR. org, 2017.
- 415 Taesup Kim, Jaesik Yoon, Ousmane Dia, Sungwoong Kim, Yoshua Bengio, and Sungjin Ahn.
416 Bayesian model-agnostic meta-learning. *arXiv preprint arXiv:1806.03836*, 2018.
- 417 Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. *International*
418 *Conference on Learning Representations*, 2016.
- 419 Andrew Gordon Wilson, Zhiting Hu, Ruslan Salakhutdinov, and Eric P Xing. Deep kernel learning.
420 In *Artificial Intelligence and Statistics*, pages 370–378, 2016.
- 421 Bernhard Scholkopf and Alexander J Smola. *Learning with kernels: support vector machines,*
422 *regularization, optimization, and beyond*. MIT press, 2001.

- 423 Ingo Steinwart and Andreas Christmann. *Support vector machines*. Springer Science & Business
424 Media, 2008.
- 425 Christopher KI Williams and Matthias Seeger. Using the nyström method to speed up kernel machines.
426 In *Advances in neural information processing systems*, pages 682–688, 2001.
- 427 Andrew Wilson and Hannes Nickisch. Kernel interpolation for scalable structured gaussian processes
428 (kiss-gp). In *International Conference on Machine Learning*, pages 1775–1784, 2015.
- 429 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz
430 Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information
431 processing systems*, pages 5998–6008, 2017.
- 432 Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Ruslan R Salakhutdinov,
433 and Alexander J Smola. Deep sets. In *Advances in neural information processing systems*, pages
434 3391–3401, 2017.
- 435 Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeswar, Sherjil Ozair, Yoshua Bengio, Aaron
436 Courville, and R Devon Hjelm. Mine: mutual information neural estimation. *arXiv preprint
437 arXiv:1801.04062*, 2018.
- 438 Bernhard Schölkopf, Sebastian Mika, Chris JC Burges, Philipp Knirsch, Klaus-Robert Müller, Gunnar
439 Rätsch, and Alexander J Smola. Input space versus feature space in kernel-based methods. *IEEE
440 transactions on neural networks*, 10(5):1000–1017, 1999.
- 441 Sambarta Dasgupta, Kumar Sricharan, and Ashok Srivastava. Finite rank deep kernel learning. 2018.
- 442 Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales.
443 Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE
444 Conference on Computer Vision and Pattern Recognition*, pages 1199–1208, 2018.
- 445 Jongmin Kim, Taesup Kim, Sungwoong Kim, and Chang D Yoo. Edge-labeling graph neural network
446 for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern
447 Recognition*, pages 11–20, 2019a.
- 448 Eleni Triantafillou, Tyler Zhu, Vincent Dumoulin, Pascal Lamblin, Kelvin Xu, Ross Goroshin, Carles
449 Gelada, Kevin Swersky, Pierre-Antoine Manzagol, and Hugo Larochelle. Meta-dataset: A dataset
450 of datasets for learning to learn from few examples. *arXiv preprint arXiv:1903.03096*, 2019.
- 451 Marta Garnelo, Jonathan Schwarz, Dan Rosenbaum, Fabio Viola, Danilo J Rezende, SM Eslami, and
452 Yee Whye Teh. Neural processes. *arXiv preprint arXiv:1807.01622*, 2018a.
- 453 Jonathan Gordon, John Bronskill, Matthias Bauer, Sebastian Nowozin, and Richard E Turner. Meta-
454 learning probabilistic inference for prediction. *arXiv preprint arXiv:1805.09921*, 2018.
- 455 Marta Garnelo, Dan Rosenbaum, Chris J Maddison, Tiago Ramalho, David Saxton, Murray Shanahan,
456 Yee Whye Teh, Danilo J Rezende, and SM Eslami. Conditional neural processes. *arXiv preprint
457 arXiv:1807.01613*, 2018b.
- 458 Hyunjik Kim, Andriy Mnih, Jonathan Schwarz, Marta Garnelo, Ali Eslami, Dan Rosenbaum, Oriol
459 Vinyals, and Yee Whye Teh. Attentive neural processes. *arXiv preprint arXiv:1901.05761*, 2019b.
- 460 Pavel G Polishchuk, Timur I Madzhidov, and Alexandre Varnek. Estimation of the size of drug-like
461 chemical space based on gdb-17 data. *Journal of computer-aided molecular design*, 27(8):675–679,
462 2013.
- 463 Gemma Roig Ngai-Man Cheung Yi Loo, Swee Kiat Lim. Few-shot regression via learned basis
464 functions. *open review preprint:r1ldYi9rOV*, 2019.
- 465 Ivan Olier, Nouredin Sadawi, G Richard Bickerton, Joaquin Vanschoren, Crina Grosan, Larisa
466 Soldatova, and Ross D King. Meta-qsar: a large-scale application of meta-learning to drug design
467 and discovery. *Machine Learning*, 107(1):285–311, 2018.

468 Appendices

469 A REGULARIZATION IMPACT

470 A.1 TASK REGULARIZATION

471 Table 6 presents the hyperparameter combinations used in the experiments to assess the impact of
 472 the trade-off parameter γ_{task} . We report the MSE performance obtained on the meta-test for each
 473 combination. To make reading this table easier, we also repeat the Fig. 5 showing the improvement
 474 of the MSE relative to $\gamma_{task} = 0$ (no regularization).

Table 6: Effect of using task regularization (parameter γ_{task}) on the MSE performance

algorithm	K	γ_{pseudo}	γ_{task} Configuration	0.00	0.01	0.10
ADKL-KRR	20	0.01	a	0.0585	0.0327	0.0289
		0.00	b	0.4051	0.2944	0.3671
		0.10	c	0.4363	0.2964	0.2882
ADKL-GP	5	0.10	d	2.4920	2.2511	2.2994
ADKL-KRR	20	0.00	e	0.0574	0.0305	0.0302
ADKL-GP	5	0.01	f	2.5611	2.1511	2.2112
		0.01	g	3.2933	2.7663	3.0971
		0.01	h	0.7675	0.7105	0.4352
ADKL-KRR	20	0.00	i	0.1201	0.0873	0.0646
		0.10	j	0.0575	0.0447	0.0273




Figure 5: Relative improvement of the MSE depending on the γ_{task} parameter


475 For a more in-depth analysis, we show below the similar tables and figures for different values of K (
 476 5, 10 and 20). These results confirm that regularizing the task encoder is helpful for any value of
 477 K , even though the impact seems to become much more important as K increases (observe that the
 478 maximum improvement in each figure increases with K).

479 **For $K = 5$**

algorithm	γ_{task} γ_{pseudo}	0.00	0.01	0.10
ADKL-GP	0.01	3.2933	2.7663	3.0971
	0.00	2.8528	3.1136	2.2801
	0.01	2.5611	2.1511	2.2112
	0.10	2.4920	2.2511	2.2994
	0.10	1.7123	1.7079	1.2808
ADKL-KRR	0.00	1.6344	1.6655	1.1974
	0.01	1.6868	1.6532	1.2173
	0.00	1.1951	1.2129	1.1998
	0.01	1.1655	1.1611	1.1416
	0.10	1.1658	1.1716	1.1442


γ_{task}	1	2	3	4	5	6	7	8	9	10
0.1	-16.0	+9.1	-16.0	-9.7	-0.3	+1.9	-2.0	+1.5	-0.4	+0.5
0.01	-6.0	-20.1	-13.7	-7.7	-25.2	-26.7	-27.8	+0.4	-2.0	-1.9

481 **For $K = 10$**



algorithm	γ_{task}	0.00	0.01	0.10
	γ_{pseudo}			
ADKL-GP	0.00	0.6423	0.6556	0.6079
	0.01	0.7675	0.7105	0.4352
	0.10	0.6182	0.6577	0.5244
ADKL-KRR	0.10	0.7326	0.6294	0.7663
	0.00	0.4051	0.2944	0.3671
	0.01	0.4386	0.3544	0.3628
	0.10	0.4363	0.2964	0.2882
	0.00	0.3170	0.2967	0.2395
	0.01	0.3070	0.2888	0.2299
	0.10	0.3038	0.2893	0.2326

483 For $K = 20$



algorithm	γ_{task}	0.00	0.01	0.10
	γ_{pseudo}			
ADKL-GP	0.00	0.1201	0.0873	0.0646
	0.01	0.0958	0.0761	0.0952
	0.10	0.0940	0.0882	0.1286
ADKL-KRR	0.01	0.1069	0.1029	0.1144
	0.00	0.0526	0.0535	0.0430
	0.01	0.0375	0.0325	0.0414
	0.10	0.0380	0.0325	0.0395
	0.00	0.0574	0.0305	0.0302
	0.01	0.0585	0.0327	0.0289
	0.10	0.0575	0.0447	0.0273

485 A.2 PSEUDO-INPUT REPRESENTATIONS

486 Table 7 presents the hyperparameter combinations used in the experiments to assess the impact of
 487 the trade-off parameter γ_{pseudo} , which governs the penalty applied to the divergence between the
 488 distribution of learned pseudo-representations and the distribution of actual representations. We also
 489 repeat in Fig. 6, the relative improvement of MSE compared to $\gamma_{pseudo} = 0$ as shown in the main
 490 text.

Table 7: Effect of the pseudo-examples regularization (parameter γ_{pseudo}) on the MSE performance

algorithm	K	γ_{task}	γ_{pseudo}	0.00	0.01	0.10
			Conf.			
ADKL-GP	10	0.10	a	0.6079	0.4352	0.5244
			b	0.0873	0.0761	0.0882
ADKL-KRR	20	0.00	c	0.0526	0.0375	0.0380
ADKL-GP	5	0.10	d	2.2801	2.2112	2.2994
ADKL-KRR	20	0.01	e	0.0535	0.0325	0.0325
ADKL-GP	5	0.01	f	2.9466	2.7663	2.7121
			g	0.1147	0.1144	0.0870
	20	0.10	h	0.1201	0.0958	0.0940
			0.00	i	3.1136	2.1511
	5	0.01	j	2.8528	2.5611	2.4920
		0.00				





Figure 6: Relative improvement of the MSE depending on the γ_{task} parameter


491 Once again, for a more in-depth analysis, we show below the same format of tables and figures for
 492 different values of K , confirming again that regularizing using the pseudo-representation can be very
 493 helpful for any value of K . It is worth noticing here that the improvement gain is more consistent for

494 $K = 5$ compared to $K \in \{10, 20\}$, supporting the fact that improving kernel maps evaluations using
 495 pseudo-representations is critical as size of the support set decreases.


496 **For $K = 5$**



498 **For $K = 10$**



500 **For $K = 20$**



502 Overall, the effect of the regularization is beneficial, even though we witness a few pathological cases.

503 A.3 JOINT IMPACT OF γ_{task} AND γ_{pseudo}

504 Since both γ_{task} and γ_{pseudo} have a high impact on the training and the generalization performance,
 505 we need to assess the relationship between the two. Fig. 7 shows, for different values of K , the
 506 relative improvement of the test MSE compared to the case where no regularization is done, i.e.
 507 $\gamma_{task} = 0$ and $\gamma_{pseudo} = 0$. Overall, one can see that higher is better in both dimensions but there
 508 seems to be a sweet spot on the grid for each value of K and therefore we can only advise the user to
 509 cross-validate on those hyperparameters.




Figure 7: Average relative improvement of the MSE and joint impact of γ_{task} and γ_{pseudo} .

510 **B PREDICTION CURVES ON THE SINUSOIDS COLLECTION**

511 Figure 8 presents a visualization of the results obtained by each model on three tasks taken randomly
 512 from the meta-test set. We provide the model with ten examples from an unseen task consisting of
 513 a slightly noisy sine function (shown in blue), and present in orange the predictions made by the
 514 network based on these ten examples.

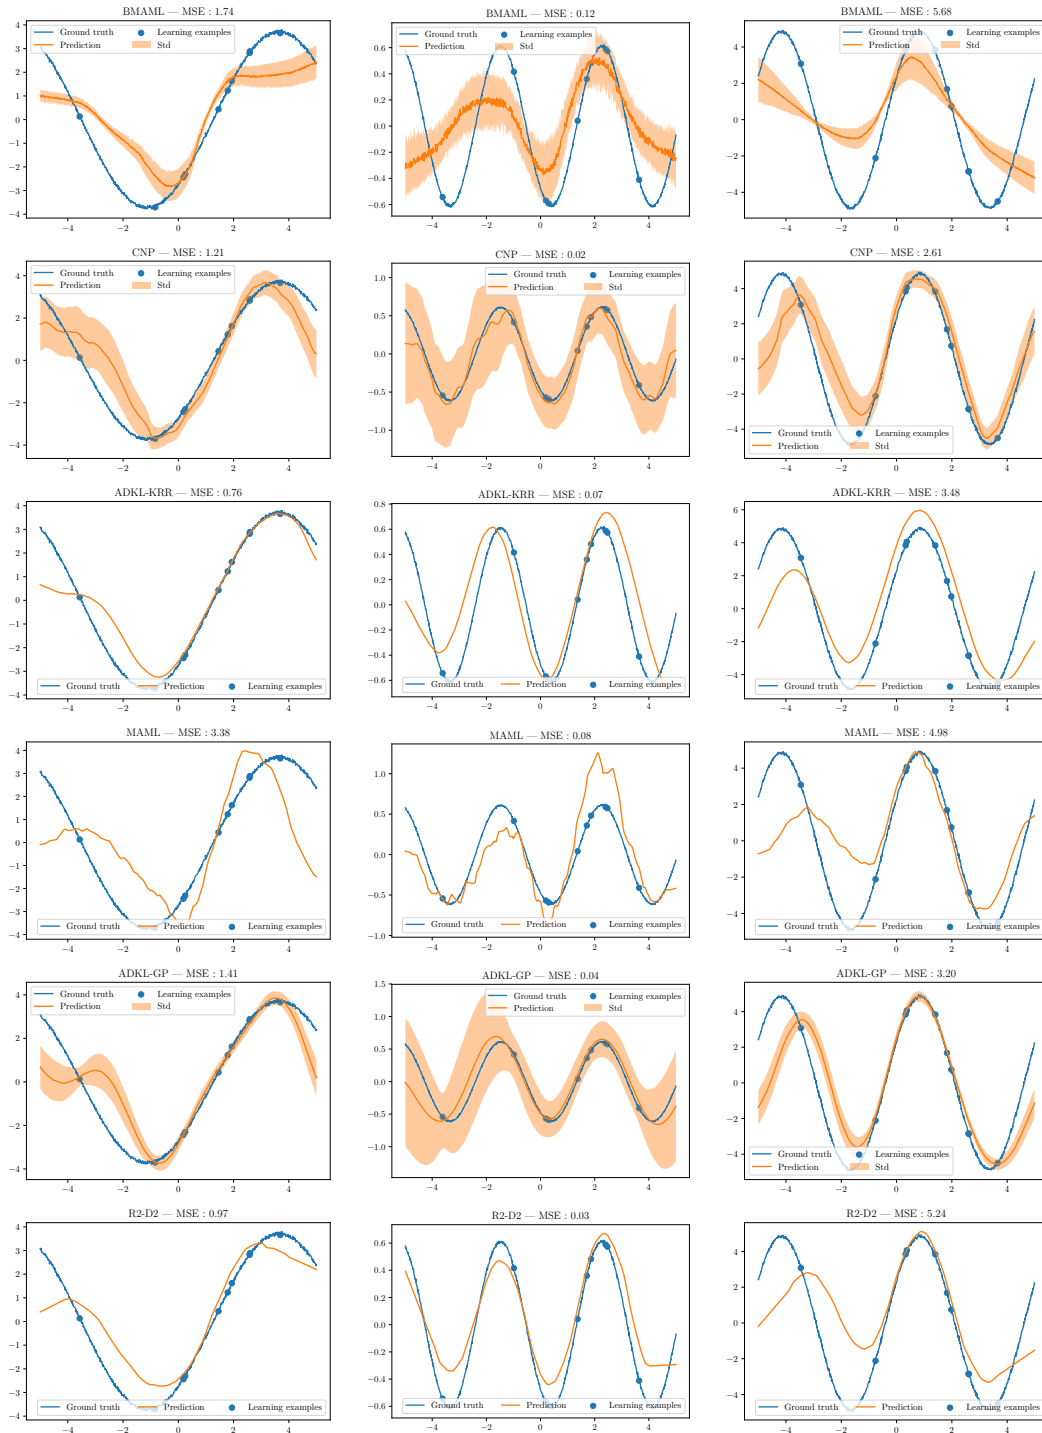


Figure 8: Meta-test time predictions on the Sinusoids collection

515 C SUPPLEMENTARY RESULTS FOR THE REAL-WORLD DATASETS

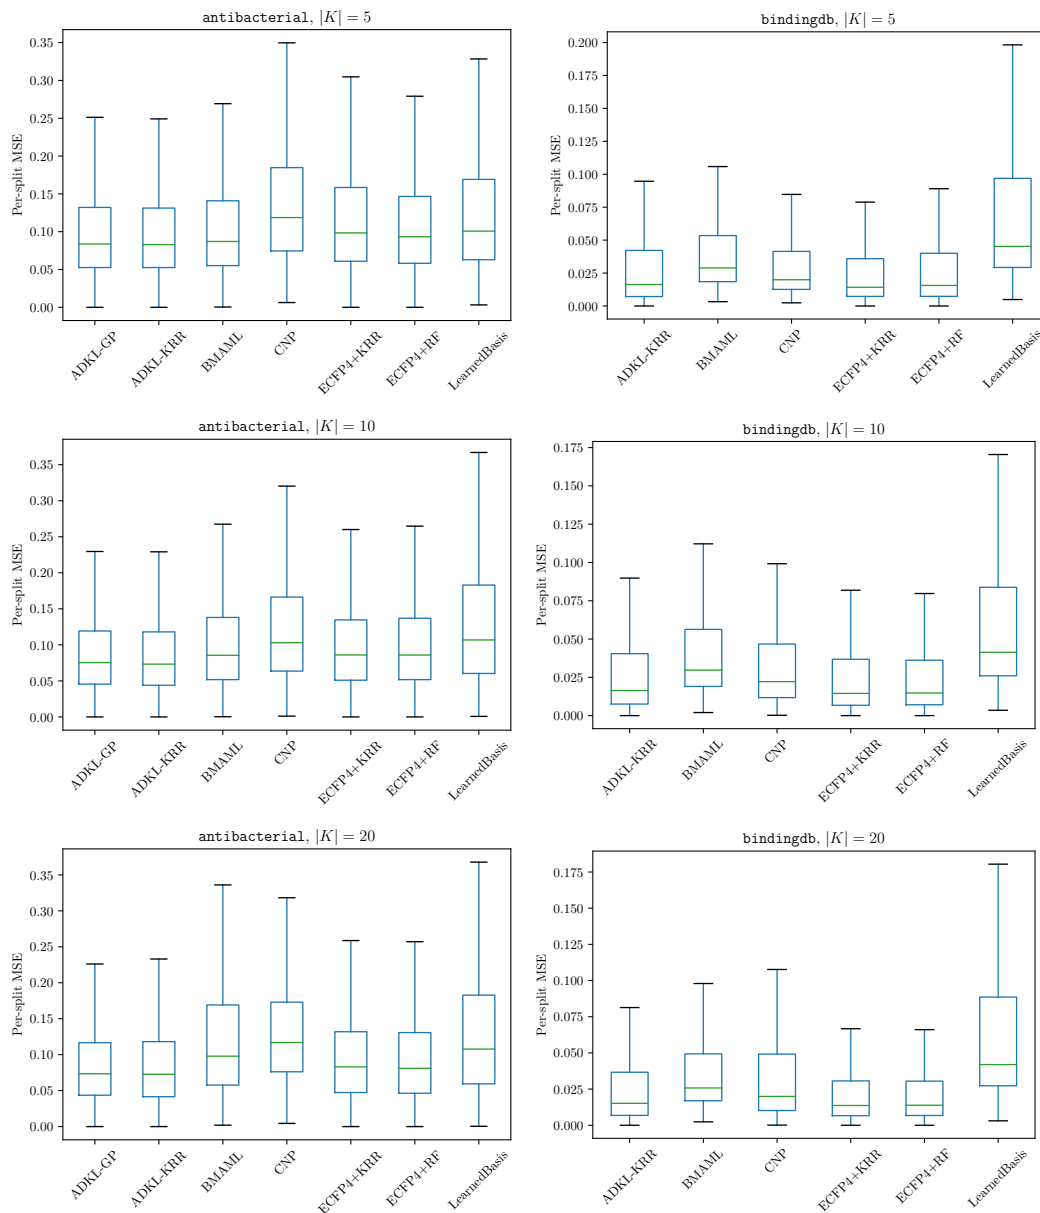


Figure 9: Distribution of the mean squared error (MSE) across the tasks

516 Figure 9 shows the distribution over the random support/query sets generated at test time. Note that
 517 the results presented in the main paper estimate the influence of the initialisation by using multiple
 518 seeds and computing the standard deviation on the *average MSE* (averaged over the support/query
 519 splits).

520 The two pieces of information are important : the results presented above give us a better idea of
 521 the “*meta-generalisation*” capabilities of each algorithm, while those in the main paper assess the
 522 reproducibility and the statistical significance of the relative improvements.