

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 one 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 a struggle 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 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 data is
38 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 differ in two aspects: the **nature of the**
 47 **meta-knowledge captured** and the **amount of adaptation performed at test-time** for new tasks or
 48 datasets. Due to the size of the total chemical space accessible when modelling bio-assays (Bohacek
 49 et al., 1996), there is a particular need for the meta-knowledge to be sufficiently rich so as to allow for
 50 extrapolation and uncertainty estimation in unseen regions of chemical space at test-time (i.e. for new
 51 tasks). Given that the same molecule can behave differently across different assays, greater test-time
 52 adaptation is also required and must be accounted for during modelling.

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

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

74 2 DEEP KERNEL LEARNING

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

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

$$h^t := \underset{h \in \mathcal{R}}{\operatorname{argmin}} \lambda \|h\|_{\mathcal{R}}^2 + \ell(h; D_{trn}^t) \quad (1)$$

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

$$h^t(\mathbf{x}) = \sum_{(\mathbf{x}_i; y_i) \in D_{trn}^t} \alpha_i k(\cdot; (\mathbf{x}_i; y_i)) \quad (2)$$

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

As DKL inherits from deep learning and kernel methods, it follows that gradient descent algorithms are required to optimize θ , which can be high dimensional such that seeing a significant amount of training samples is essential to avoid overfitting. However, once the latter condition is met, scalability of the kernel method becomes limiting as the running time of kernel methods scales approximately in $O(m^3)$ for a training set of m samples. Some approximations of the kernel are thus needed for the scalability of the DKL method (see Williams and Seeger (2001); Wilson and Nickisch (2015)).

Few-Shot DKL: In the setup of episodic meta learning, also known as few-shot learning, one has access to a meta-training collection $\mathcal{D}_{meta-trn} := \left\{ (D_{trn}^t; D_{val}^t) \right\}_{t=1}^T$ of T tasks to learn how to learn from few datapoints. Each task t_j has its own training (or support) set D_{trn}^t and validation (or query) set D_{val}^t . A meta-testing collection $\mathcal{D}_{meta-tst}$ is also available to assess the generalization performance of the few-shot algorithm across unseen tasks. To obtain a Few-Shot DKL (FSDKL) method for FSR in such settings, one can share the parameters θ across all tasks, similar to metric learning algorithms. Hence, for a given task t_j , the inputs are first transformed by the function ϕ and then a kernel method is used to obtain the regressor h^t , which will be evaluated on D_{val}^t . Here, KRR and GP are explored as they are the state-of-the-art algorithms for kernel-based regression. The latter is used to allow our models to provide accurate predictive uncertainty, which is useful when modelling biological assays.

KRR: Using the squared loss and the L2-norm to compute khk_R , KRR gives the optimal regressor for a task t and its validation loss $L^t_{\theta; \theta}$ as follows:

$$h^t(\mathbf{x}) = K_{\mathbf{x};trn} \left(K_{trn;trn} + I \right)^{-1} \mathbf{y}_{trn} \tag{3}$$

$$L^t_{\theta; \theta} = \mathbf{E}_{\mathbf{x};y} \mathbf{E}_{D_{val}^t} \left(K_{\mathbf{x};trn}(\mathbf{x}; \mathbf{y}) \right)^2 \tag{4}$$

where $\mathbf{y}_{trn} = (y_1; \dots; y_{|D_{trn}^t|})^T$, $K_{trn;trn}$ is the matrix of kernel evaluations and each entry is $k(\langle \mathbf{x}_i; \mathbf{x}_j \rangle; \theta)$ for $(\mathbf{x}_i; \mathbf{x}_j) \in D_{trn}^t \times D_{trn}^t$. An equivalent definition applies to $K_{\mathbf{x};trn}$.

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

$$L^t_{\theta; \theta} = -\ln \mathcal{N}(\mathbf{y}_{val}; \mathbf{E}[h^t]; \text{cov}(h^t)); \tag{5}$$

$$\mathbf{E}[h^t] = K_{val;trn} \left(K_{trn;trn} + I \right)^{-1} \mathbf{y}_{trn}; \tag{6}$$

$$\text{cov}(h^t) = K_{val;val} - K_{val;trn} \left(K_{trn;trn} + I \right)^{-1} K_{trn;val} \tag{7}$$

Finally, the parameters θ of the neural network, along with ϕ and the kernel hyperparameters θ_k , are optimized using the expected loss on all tasks:

$$\text{argmin}_{\theta; \theta_k} \mathbf{E}_{t \sim \mathcal{D}_{meta-trn}} L^t_{\theta; \theta} \tag{8}$$

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

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

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

127 3 ADAPTIVE DEEP KERNEL LEARNING

128 FSDKL uses a shared deep kernel function, in which the base kernel k is user-chosen (although its
 129 parameters are learned during meta-training). Given that the choice of the kernel function dictates
 130 almost all the generalization properties of any kernel-based method, it may not be optimal to leave it
 131 to the user. In addition, for modelling bio-assays, it is not straightforward how to best incorporate
 132 task-specific prior knowledge in this shared and user-chosen kernel. We overcome these limitations
 133 by developing the Adaptive Deep Kernel Learning (ADKL) framework, illustrated by Fig. 1.

134 ADKL also aims to obtain a non-linear embedding of inputs using a deep neural network shared
 135 by all tasks before finding the minimal norm task-specific regressor h^t using either GP or KRR
 136 as described in Section 2 (ADKL-GP and ADKL-KRR will refer to our algorithm when using
 137 GP and KRR, respectively). The fundamental difference between FSDKL and ADKL lies in the
 138 kernel definition, which brings significantly more flexibility in the latter case relative to the former.
 139 Specifically, during the meta-training, ADKL *learns to learn* task-specific kernel functions instead of
 140 using one chosen by the user. It does so by learning how to represent tasks with the task encoding
 141 network and then how to leverage task embeddings to build task-specific kernels using a multi-
 142 modal neural network C . Given a task t , ADKL thus first computes its embedding $\mathbf{z}_t = \psi_\eta(D_{trn}^t)$
 143 using its support set D_{trn}^t and deduces the adapted kernel with C . We describe in more detail both
 144 the task encoding network and the network C responsible for computing the task-specific kernel
 145 below.

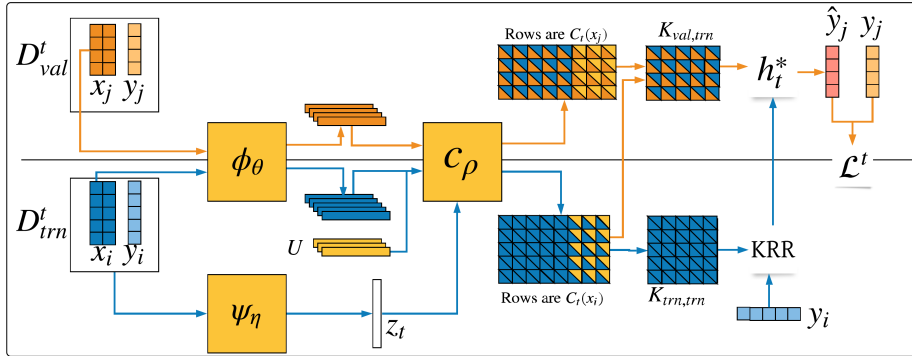


Figure 1: ADKL-KRR. The blue and orange colors show the procedure for a task during internal train and test, respectively.

146 3.1 TASK ENCODING

147 The challenge of the network is to capture complex dependencies in the training set D_{trn}^t to
 148 provide a useful task encoding \mathbf{z} . Furthermore, the task encoder should be invariant to permutations
 149 of the training set and be able to encode a variable amount of samples. After exploring a variety
 150 of architectures, we found that those that are more complex, such as Transformers (Vaswani et al.,
 151 2017), tend to underperform. This is possibly due to overfitting or the sensitivity of training such
 152 architectures.

153 Consequently, inspired by DeepSets (Zaheer et al., 2017), we propose the following order in-
 154 variant network. It begins its computations by representing each input-target pair $\mathbf{xy}_i =$
 155 $\mathbf{r}(\text{Concat}(\mathbf{x}_i; \mathbf{v}(y_i)))$ for all $(\mathbf{x}_i; y_i) \in D_{trn}^t$, using neural networks \mathbf{v} and \mathbf{r} . The \mathbf{xy}_i
 156 captures nonlinear interactions between the inputs and the targets if \mathbf{r} is a nonlinear transforma-
 157 tion. Then, by computing $\bar{\mathbf{x}}_{\mathbf{xy}}^t$ and $\hat{\mathbf{x}}_{\mathbf{xy}}^t$, the empirical mean and standard deviation of the set
 158 $\{\mathbf{xy}_1; \mathbf{xy}_2; \dots; \mathbf{xy}_m\}$, respectively, we obtain the task representation as follows:

$$\mathbf{z}_t = \psi_\eta(D_{trn}^t) := \left[\bar{\mathbf{x}}_{\mathbf{xy}}^t; \hat{\mathbf{x}}_{\mathbf{xy}}^t \right] \quad (9)$$

159 As $\bar{\mathbf{x}}_{\mathbf{xy}}^t$ and $\hat{\mathbf{x}}_{\mathbf{xy}}^t$ are invariant to permutations in D_{trn}^t , it follows that \mathbf{z}_t is also permutation invariant.
 160 Overall, \mathbf{z}_t is simply the concatenation of the first and second moments of the sample representations,

161 which were nonlinear transformations of the original inputs and targets. The learnable parameters
162 of the task encoder include all the parameters of the networks \mathbf{v} and \mathbf{r} , and are shared across all tasks.

163 To help the training of these parameters, we maximize the mutual information between D_{trn}^t and
164 D_{val}^t i.e. we expect and encourage the network to produce similar task encodings using any data
165 partitions for a given task. More explicitly, we use the MINE algorithm (Belghazi et al., 2018), which
166 optimizes a lower bound on the mutual information. For two random variables $r; s \sim p(r; s)$ and a
167 similarity measure f between r and s , parameterized by θ , the following inequality holds:

$$I[r; s] \leq \max_{\theta} \mathbf{E}_{r; s} f_{\theta}(r; s) - \ln \mathbf{E}_r \mathbf{E}_s \mathbf{E}_{p(s)} e^{f_{\theta}(r; s)}; \quad (10)$$

Using a batch of b tasks², and the cosine similarity c as the similarity measure between two task encodings, one obtains and maximizes $I[D_{trn}; D_{val}] \approx t$, where:

$$t \stackrel{\text{def}}{=} \frac{1}{b} \sum_{j=1}^b c(D_{trn}^j; D_{val}^j) - \ln \frac{1}{b(b-1)} \sum_{j=1}^b \sum_{i \neq j} e^{c(D_{trn}^j; D_{val}^i)}; \quad (11)$$

168 3.2 TASK-SPECIFIC KERNEL

169 Let H and Z be the output domains of \mathbf{v} and \mathbf{r} respectively. We define a pairwise function on H ,
170 whose outputs are dependant from the task representations in Z as follows:

$$c : H \times H \times Z \rightarrow \mathbb{R} \\ c(\mathbf{x}; \mathbf{x}^0; \mathbf{z}_t) = \text{MLP}([\|\mathbf{x} - \mathbf{x}^0\|; \mathbf{z}_t]); \quad (12)$$

171 where $[\cdot; \cdot]$ is the concatenation operator. It bears mentioning that the parameters θ are shared across
172 all tasks and learned during the meta-training. Also, c is symmetric and stationary with regard to
173 its inputs (\mathbf{x} and \mathbf{x}^0) as their element-wise L2 distances vector is received as input of the fully
174 connected network. Further, by simply concatenating the task representation \mathbf{z}_t to this distance
175 vector at the input, c provides a powerful approach to producing task-specific kernels. However,
176 these kernels are not positive semi-definite (PSD) and cannot be directly used for KRR and GP.
177 Therefore, without losing any information given by c , we compute task-specific PSD kernels $k_{\cdot; t}$
178 as the empirical kernel maps with regard to the support set inputs, i.e.:

$$k_{\cdot; t}(\mathbf{x}; \mathbf{x}^0) = C_t(\mathbf{x}) C_t(\mathbf{x}^0); \quad \text{with} \\ C_t(\mathbf{x}) = (c(\mathbf{x}; \mathbf{x}_1; \mathbf{z}_t); \dots; c(\mathbf{x}; \mathbf{x}_m; \mathbf{z}_t)); \quad \text{and } (\mathbf{x}_i;) \in D_{trn}^t \delta_i = 1; \quad ; m \quad (13)$$

179 Using the empirical kernel map of c to compute $k_{\cdot; t}$ offers the opportunity to improve the kernel
180 evaluations in low data settings using some unlabelled data. More precisely, instead of computing the
181 empirical kernel map with regard to only D_{trn}^t , we could use $(D_{trn}^t \cup U)$ where U is a set of unlabelled
182 inputs. However, to avoid a significant increase in the computation costs of the PSD kernels, $|U|$
183 should be kept relatively small (in our experiments $|U| \leq 50$). One must also be careful about the
184 composition of U to avoid overfitting certain tasks and under-fitting others. Therefore, instead of
185 asking the user to provide the set U , we propose directly learning them through back-propagation. To
186 do so, we introduce *pseudo-input representations* (or *pseudo-representations*) $\mathbf{u}_l \in H$ that are shared
187 by all tasks and learned during meta-training. The function C_t , from Eq. (13), becomes:

$$C_t(\mathbf{x}) = (c(\mathbf{x}; \mathbf{x}_1; \mathbf{z}_t); \dots; c(\mathbf{x}; \mathbf{x}_m; \mathbf{z}_t); c(\mathbf{x}; \mathbf{u}_1; \mathbf{z}_t); \dots; c(\mathbf{x}; \mathbf{u}_l; \mathbf{z}_t)); \quad (14) \\ \text{with } \mathbf{u}_l \in U \delta_l = 1; \quad ; l \in [1; l] \text{ and } (\mathbf{x}_i;) \in D_{trn}^t \delta_i = 1; \quad ; m$$

188 These *pseudo-representations* can be thought of as parameters of the adaptive kernel and the number
189 to be included is a hyperparameter of the algorithm. To prevent their collapse into a single point
190 and ensure that they are well distributed in the feature space H , we add a regularization term to the
191 training loss. Let p and q be the distributions that generate the true input representations and the
192 pseudo-input representations, respectively. We make the assumption that p and q are both multivariate
193 Gaussian distributions with diagonal covariance matrices and have respective parameters (μ, Σ)

²This yields a small bias on the gradient since the right hand side takes the log of the expectations. Since we are not interested in the precise value of the mutual information, this does not constitute a problem.

194 and $(\mathbf{u}, \mathbf{u}^2)$. The parameters of p are estimated using the running means and variances of all input
 195 representations computed over batches of tasks. Those of q are estimated using U . The training of
 196 the pseudo-input representations is then regularized by minimizing the KL distance $\mathcal{D}_{\mathbf{u}}$ between p
 197 and q , i.e.:

$$\mathcal{D}_{\mathbf{u}} = KL(N(\mathbf{u}; \mathbf{u}^2) \parallel N(\mathbf{u}; \mathbf{u}^2)) \quad (15)$$

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

$$\operatorname{argmin}_{\mathbf{u}; \mathbf{u}^2; t_j} \mathbf{E}_{\mathcal{B}} L_{\mathcal{B}}^j; \dots; \mathbf{u}; \dots; \text{task } t_j + \text{pseudo } \mathcal{D}_{\mathbf{u}}; \dots \quad (16)$$

198 with task 0 as a tradeoff hyperparameter for the regularization of the task-encoder, pseudo 0
 199 as a tradeoff hyperparameter for the regularization of the pseudo-inputs.

200 4 RELATED WORK

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

212 On the spectrum between NNs and kernel methods we must also mention metric learning. Metric
 213 learning algorithms learn an input covariance function shared across tasks but rely only on the
 214 expressive power of DNNs. First, stochastic kernels are built out of shared feature extractors and
 215 simple pairwise metrics (e.g. cosine similarity (Vinyals et al., 2016), Euclidean distance (Snell et al.,
 216 2017)), or parametric functions (e.g. relation modules (Sung et al., 2018), graph neural networks
 217 (Garcia and Bruna, 2017; Kim et al., 2019a)). Then, within tasks, the predictions are distance-
 218 weighted combinations of the training labels with the stochastic kernel evaluations—no adaptation is
 219 done.

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

232 Since ADKL-GP learns task-specific stochastic processes, it is related to neural processes (Garnelo
 233 et al., 2018a) and the ML-PIP framework (Gordon et al., 2018). Both propose a scalable alternative
 234 to learning regression functions by performing inference on stochastic processes. In these families
 235 of methods, both Conditional Neural Processes (CNP) (Garnelo et al., 2018b) and Attentive Neural
 236 Processes (ANP) (Kim et al., 2019b) learn conditional stochastic processes parameterized by task-
 237 specific conditions derived from the support sets, but CNP is the most related to ADKL-GP. CNP is
 238 an instance of ML-PIP when the task encoder gives a point estimate of the task parameters instead
 239 of a distribution. Finally, the main differences between ANP and CNP are the architecture of the
 240 task-encoder and the lack of mathematical guarantees associated with stochastic processes in CNP
 241 (as it does not impose any consistency with respect to a prior process). By comparison, ADKL-GP

242 also learns conditional stochastic processes but has mathematical guarantees thanks to GP and PSD
 243 kernels.

244 5 DATASETS

245 Existing FSR methods have been mostly tested on 1D function regression and pixel-wise image
 246 completion tasks with MNIST and CelebA (Kim et al., 2018; Garnelo et al., 2018b;a). On one hand,
 247 the 1D regression tasks are all relatively simple, almost noise-less, and homogeneous. On the other
 248 hand, methods have been successful for image completion tasks only outside the few-shot regime (i.e.
 249 when the number of samples is greater than 500) (Garnelo et al., 2018b;a). For these reasons, we
 250 introduce two task collections from a real-world context. Deemed **Binding** and **Antibacterial**, these
 251 task collections contain data from bio-assays that are representative of real-world FSR tasks in drug
 252 discovery. The pre-processed versions of these collections and detailed statistics are available [here](#)
 253 ([anonymized link](#)).

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

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

265 Their meta-test partitions each contain 500 tasks, with the remaining used in the meta-train and
 266 meta-validation. The molecules (represented as **SMILES**) are converted into vectors using routines
 267 available in the RDKit software (more precisely into **ECFP6** binary fingerprint vectors of 4,096
 268 dimensions). These inputs were also processed in all methods using the same feature extractor
 269 architecture, which is a fully-connected network of 256 256 256. Due to the high noise-to-signal
 270 ratio, the targets are first *log2*-scaled and then scaled linearly between 0 and 1 to avoid scaling issues
 271 during training.

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

286 To test our method in a noise-less environment, we also use the **Sinusoids** collection introduced by
 287 Kim et al. (2018). This challenging few-shot regression benchmark consists of 5,000 tasks defined
 288 by functions of the form: $y = A \sin(wx + b) +$ with $A \geq [0.1; 5.0]$, $b \geq [0.0; 2]$, and
 289 $w \geq [0.5; 2.0]$. Sampling inputs $x \geq [5.0; 5.0]$ and observational noise $\geq N(0; (0.01A)^2)$ and
 290 computing y gives the samples for each task. Here, the meta-train, meta-validation, and meta-test
 291 contain 56.25%, 18.75% and 25% of all the tasks, respectively, and all methods use the same feature
 292 extractor architecture, which is a fully-connected network of 40 40 40.




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.

293 6 EXPERIMENTS

294 6.1 BENCHMARKING ANALYSIS

295 Performance of ADKL is evaluated against a FSDKL instance (R2-D2 of Bertinetto et al. (2018)),
 296 CNP (Garnelo et al., 2018b), MAML (Finn et al., 2017), BMAML (Kim et al., 2018), ProtoMAML
 297 (Triantafillou et al., 2019) and Learned Basis (Yi Loo, 2019) (all implementations are available
 298 [here \(anonymized link\)](#)). These algorithms have all proven to have efficient and effective test-time
 299 adaptation routines and therefore constitute strong baselines for benchmarking. Tables 1 to 3 report
 300 the average MSE over all tasks and 20 random support and query partitions for each task, for different
 301 sized support sets.

302 For the Sinusoids collection, we observe that DKL-based methods significantly outperform all
 303 other methods despite their test-time adaptation capabilities. These results alone demonstrate the
 304 effectiveness of DKL-based methods in FSR relative to the current state-of-the-art. Furthermore, of
 305 all DKL-based methods, ADKL-KRR shows consistently stronger performance than others. This
 306 demonstrates that using ADKL increases test-time performance relative to FS-DKL (as R2-D2 and
 307 ADKL-KRR only differ by the kernel definition). It also indicates that attempting to capture the
 308 model uncertainty using GP in ADKL (instead of KRR) comes with a significant cost, especially in
 309 lower data regimes. This may be due to the inability of GP to differentiate between the observational
 310 noise and the model uncertainty as the number of samples get smaller. It is also important to notice
 311 that all methods using the task representation significantly outperform those that do not. This shows
 312 that adequately capturing the task representation is crucial for this task collection, which ADKL-KRR
 313 appears to be well-equipped to handle.

314 Tables 2 and 3 show that real-world datasets are challenging for most methods, as their MSE only
 315 marginally improves when the size of the support set increases. It should be noted that while the
 316 scaling of the targets makes the MSE low for all methods, even a decrease of 0.005 can translate
 317 into large improvements of the modelling accuracy. However, we still observe that ADKL-KRR
 318 outperforms all other methods when the number of samples is greater than 10, again providing
 319 evidence of effectiveness of our method for FSR with complex task distributions. The gaps between
 320 ADKL and R2-D2 for these collections also confirm that using task specific kernels can be very
 321 useful even though inferring the right kernel can become difficult as the size of the support set gets
 322 smaller. Finally, it also appears that estimating the model uncertainty using ADKL-GP instead of
 323 ADKL-KRR comes with a marginal accuracy cost.

324 6.2 ACTIVE LEARNING

325 In this section, we report the results of active learning experiments. Our intent is to measure the
 326 effectiveness of the uncertainty captured by the predictive distribution of ADKL-GP for active

m	5	10	20
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
ProtoMAML	2.044	1.369	0.846
FSDKL(R2D2)	1.002	0.073	0.009

Table 1: Average MSE on Sinusoidals

m	5	10	20
BMAML	0.061	0.059	0.057
CNP	0.064	0.062	0.061
Learned Basis	0.063	0.060	0.059
MAML	-	-	-
ADKL-GP	0.064	0.056	0.051
ADKL-KRR	0.063	0.054	0.051
ProtoMAML	0.061	0.059	0.065
FSDKL(R2D2)	0.060	0.060	0.055

Table 2: Average MSE on Binding

m	5	10	20
BMAML	0.067	0.060	0.060
CNP	0.070	0.069	0.068
Learned Basis	0.068	0.065	0.093
MAML	-	-	-
ADKL-GP	0.068	0.064	0.060
ADKL-KRR	0.068	0.059	0.058
ProtoMAML	0.065	0.063	0.070
FSDKL(R2D2)	0.066	0.064	0.063

Table 3: Average MSE on Antibacterial

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

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

6.3 ABLATION EXPERIMENTS

In our final set of experiments, we more closely evaluate the impact of the task encoder and the pseudo-inputs on the generalization during meta-testing. We do so by training and evaluating ADKL on Sinusoids with different hyperparameter combinations. Figs. 4a to 4d show the relative improvements (negative values) or setbacks (positive values) in the meta-test MSE compared to different baselines (but the joint impact of $task$ and $pseudo$ is only discussed in Appendix A.3).

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

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

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

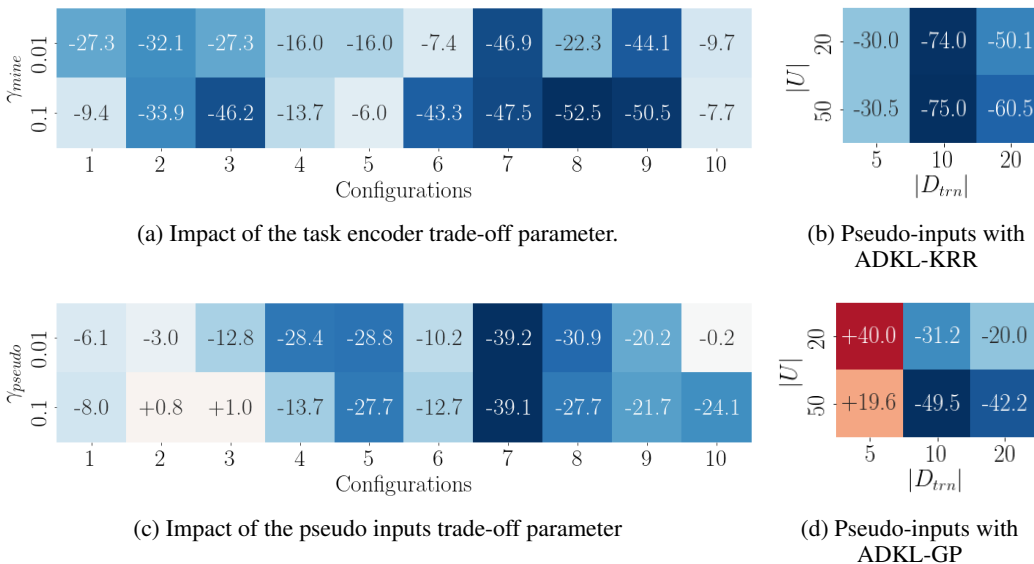


Figure 4: Relative decrease/increase in the meta-test MSE compared to different baselines. In (a) and (c) the baselines are $\gamma_{task} = 0$ and $\gamma_{pseudo} = 0$, respectively. In (b) and (d) the baselines are $jUj = 0$

375 **7 CONCLUSION**

376 In this work, we investigate the modelling of biological assays using few-shot learning methods.
 377 We propose a new framework, ADKL, that stores meta-knowledge in kernel functions and adapts
 378 to new tasks using KRR or GP. Our experiments provide evidence that the additional adaptation
 379 capacity at test-time provided by these methods increases generalization when modelling bio-assays
 380 and on 1D sinusoidal regression tasks. In a Bayesian setup, they better estimate predictive uncertainty,
 381 increasing their utility in real-world applications such as drug discovery. Finally, by making our
 382 bio-assay task collections publicly available, we hope that the community will leverage them to
 383 propose FSR algorithms that are ready to be deployed under real-world constraints, with the ultimate
 384 aim of accurately predicting key molecular properties early in the drug discovery pipeline.

385 REFERENCES

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

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

468

Appendices

469

A REGULARIZATION IMPACT


470

A.1 TASK REGULARIZATION

471 Table 4 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 4: 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	1	0.0585	0.0327	0.0289
		0.00	2	0.4051	0.2944	0.3671
		0.10	3	0.4363	0.2964	0.2882
ADKL-GP	5	0.10	4	2.4920	2.2511	2.2994
ADKL-KRR	20	0.00	5	0.0574	0.0305	0.0302
ADKL-GP	5	0.01	6	2.5611	2.1511	2.2112
		0.01	7	3.2933	2.7663	3.0971
	10	0.01	8	0.7675	0.7105	0.4352
		0.00	9	0.1201	0.0873	0.0646
		0.10	10	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_{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
	0.00	1.7123	1.7079	1.2808
ADKL-KRR	0.01	1.6344	1.6655	1.1974
	0.10	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
0.01	-16.0	+9.1	-16.0	-9.7	-0.3	+1.9	-2.0	+1.5	-0.4	+0.5
0.1	-6.0	-20.1	-13.7	-7.7	-25.2	-26.7	-27.8	+0.4	-2.0	-1.9

481 **For $K = 10$**



483 For $K = 20$



485 A.2 PSEUDO-INPUT REPRESENTATIONS

486 Table 5 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 5: Effect of the pseudo-examples regularization (parameter γ_{pseudo}) on the MSE performance

algorithm	K	task	γ_{pseudo} Conf.	0.00	0.01	0.10
ADKL-GP	10	0.10	1	0.6079	0.4352	0.5244
			2	0.0873	0.0761	0.0882
ADKL-KRR	20	0.00	3	0.0526	0.0375	0.0380
ADKL-GP	5	0.10	4	2.2801	2.2112	2.2994
ADKL-KRR	20	0.01	5	0.0535	0.0325	0.0325
ADKL-GP	5	0.01	6	2.9466	2.7663	2.7121
			20	0.10	7	0.1147
	5	0.00	8	0.1201	0.0958	0.0940
			0.01	9	3.1136	2.1511
	5	0.00	10	2.8528	2.5611	2.4920





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 $K \geq 10$, supporting the fact that improving kernel maps evaluations using
 495 pseudo-representations is critical as size of the support set decreases.

496 **For $K = 5$**




algorithm	γ_{pseudo}	0.00	0.01	0.10
ADKL-GP	0.01	2.9466	2.7663	2.7121
	0.10	2.2801	2.2112	2.2994
ADKL-KRR	0.00	1.7123	1.6344	1.6868
	0.00	1.1951	1.1655	1.1658
ADKL-GP	0.00	2.8528	2.5611	2.4920
	0.10	1.1998	1.1416	1.1442
ADKL-KRR	0.01	1.2129	1.1611	1.1716
	0.10	1.2808	1.1974	1.2173
ADKL-GP	0.01	3.1136	2.1511	2.2511
ADKL-KRR	0.01	1.7079	1.6655	1.6532

498 **For $K = 10$**



algorithm	γ_{pseudo}	0.00	0.01	0.10
ADKL-GP	0.01	0.7329	0.7907	0.6294
	0.10	0.7479	0.7800	0.7663
ADKL-KRR	0.00	0.3170	0.3070	0.3038
	0.00	0.6423	0.7675	0.6182
ADKL-GP	0.10	0.3671	0.3628	0.2882
	0.01	0.2967	0.2888	0.2893
ADKL-KRR	0.01	0.6556	0.7105	0.6577
	0.00	0.7145	0.6758	0.7326
ADKL-GP	0.10	0.6079	0.4352	0.5244
	0.10	0.2395	0.2299	0.2326

500 **For $K = 20$**



algorithm	γ_{pseudo}	0.00	0.01	0.10
ADKL-GP	0.00	0.1201	0.0958	0.0940
	0.00	0.0794	0.1069	0.0702
ADKL-KRR	0.01	0.0873	0.0761	0.0882
	0.01	0.0305	0.0327	0.0447
ADKL-GP	0.00	0.0526	0.0375	0.0380
	0.10	0.0302	0.0289	0.0273
ADKL-KRR	0.00	0.0574	0.0585	0.0575
	0.10	0.1147	0.1144	0.0870
ADKL-GP	0.01	0.0535	0.0325	0.0325
	0.10	0.0430	0.0414	0.0395

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.

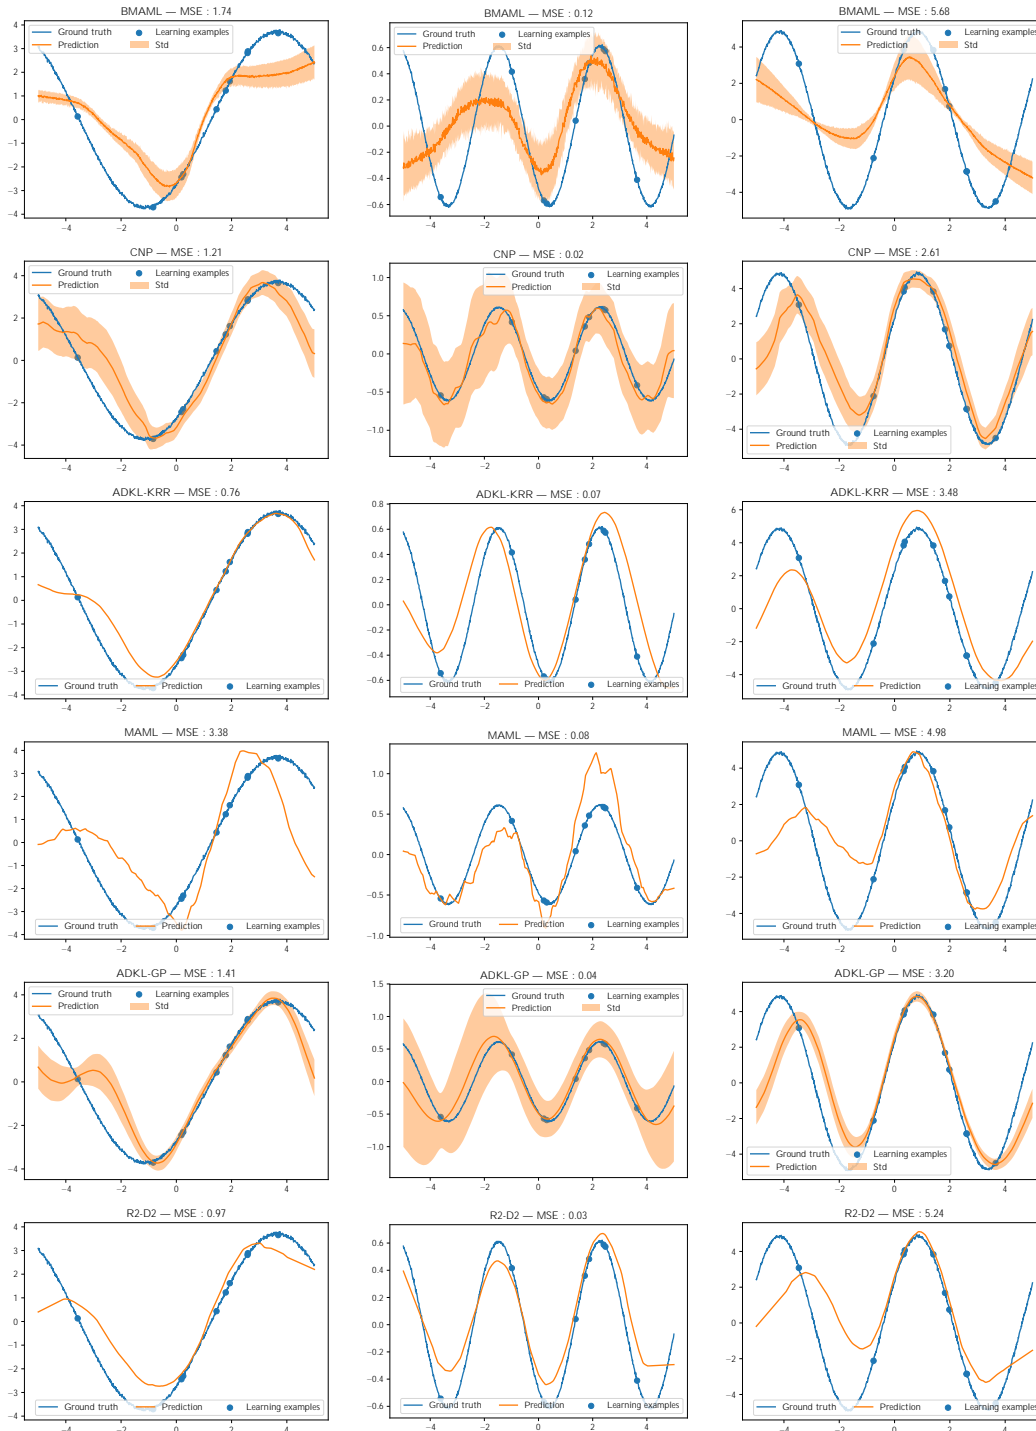


Figure 8: Meta-test time predictions on the Sinusoids collection