

LEARNING TO EXTRAPOLATE: A TRANSDUCTIVE APPROACH

Aviv Netanyahu^{1,2*}, Abhishek Gupta^{1,2,3*}, Max Simchowitz², Kaiqing Zhang^{2,4} & Pulkit Agrawal^{1,2}
 Improbable AI Lab¹ MIT² University of Washington³ University of Maryland, College Park⁴

ABSTRACT

Machine learning systems, especially with overparameterized deep neural networks, can generalize to novel test instances drawn from the same distribution as the training data. However, they fare poorly when evaluated on *out-of-support* test points. In this work, we tackle the problem of developing machine learning systems that retain the power of overparameterized function approximators while enabling extrapolation to out-of-support test points when possible. This is accomplished by noting that under certain conditions, a “transductive” reparameterization can convert an out-of-support extrapolation problem into a problem of within-support combinatorial generalization. We propose a simple strategy based on bilinear embeddings to enable this type of combinatorial generalization, thereby addressing the out-of-support extrapolation problem under certain conditions. We instantiate a simple, practical algorithm applicable to various supervised learning and imitation learning tasks.

1 INTRODUCTION

Generalization is a central problem in machine learning. Typically, one expects generalization when the test data is sampled from *the same distribution* as the training set, i.e. *out-of-sample* generalization. However, in many scenarios, test data is sampled from a different distribution from the training set, i.e., *out-of-distribution* (OOD). In some OOD scenarios, the test distribution is assumed to be known during training – a common assumption made by meta-learning methods (Finn et al., 2017b). Several works have tackled a more general scenario of “reweighted” distribution shift (Koh et al., 2021; Quinero-Candela et al., 2008) where the test distribution shares support with the training distribution, but has a different and unknown probability density. This setting can be tackled via distributional robustness approaches (Sinha et al., 2018; Rahimian & Mehrotra, 2019).

Our paper aims to find structural conditions under which generalization to test data with support *outside* of the training distribution is possible. Formally, assume the problem of learning function $h: \hat{y} = h(x)$ using data $\{(x_i, y_i)\}_{i=1}^N \sim \mathcal{D}_{\text{train}}$, where $x_i \in \mathcal{X}_{\text{train}}$, the training domain. We are interested in making accurate predictions $h(x)$ for $x \notin \mathcal{X}_{\text{train}}$ (see examples in Fig 1). Consider an example task of predicting actions to reach a desired goal (Fig 1b). During training, goals are provided from the blue cuboid ($x \in \mathcal{X}_{\text{train}}$), but test time goals are from the orange cuboid ($x \notin \mathcal{X}_{\text{train}}$). If h is modeled using a deep neural network, its predictions on test goals in the blue area are likely to be accurate, but for the goals in the orange area the performance can be arbitrarily poor unless further domain knowledge is incorporated. This challenge manifests itself in a variety of real-world problems, ranging from supervised learning problems like object classification (Barbu et al., 2019), sequential decision making with reinforcement learning (Kirk et al., 2021), transferring reinforcement learning policies from simulation to real-world (Zhao et al., 2020), imitation learning (de Haan et al., 2019), etc. Reliably deploying learning algorithms in unconstrained environments requires accounting for such “out-of-support” distribution shift, which we refer to as *extrapolation*.

It is widely accepted that if one can identify some *structure* in the training data that constrains the behavior of optimal predictors on novel data, then extrapolation may become possible. Several methods can extrapolate if the nature of distribution shift is known apriori: convolution neural

*denotes equal contribution. Correspondence to Aviv Netanyahu <avivn@mit.edu>, Abhishek Gupta <abhgupta@cs.washington.edu>.




Figure 1: In the real-world the test distribution (orange) often has a *different support* than the training distribution (blue). Some illustrative tasks: (a) grasp point prediction for object instances with out-of-support scale, (b) action prediction for reaching out-of-support goals, (c) function value prediction for an out-of-support input range. The black crosses show predictions for a conventionally trained deep neural network that makes accurate predictions for in-support inputs, but fails on out-of-support inputs. We propose an algorithm that makes accurate out-of-support predictions under a set of assumptions.

networks are appropriate if a training pattern appears at an out-of-distribution translation in a test example (Kayhan & Gemert, 2020). Similarly, accurate predictions can be made for object point clouds in out-of-support orientations by *building in* SE(3) equivariance (Deng et al., 2021b). Another way to extrapolate is if the model class is known apriori: fitting a linear function to a linear problem will extrapolate. Similarly, methods like NeRF (Zhang et al., 2022a) use physics of image formation to learn a 3D model of a scene which can synthesize images from novel viewpoints.

We propose an alternative structural condition under which extrapolation is feasible. Typical machine learning approaches are *inductive*: decision-making rules are inferred from training data and employed for test-time predictions. An alternative to induction is *transduction* (Gammerman et al., 1998) where a test example is compared with the training examples to make a prediction. Our main insight is that in the transductive view of machine learning, extrapolation can be reparameterized as a combinatorial generalization problem, which, under certain low-rank and coverage conditions (Shah et al., 2020; Agarwal et al., 2021b; Andreas et al., 2016; Andreas, 2019), admits a solution.

First we show how we can (i) reparameterize out-of-support inputs $h(x_{\text{test}}) \rightarrow h(\Delta x, x')$, where $x' \in \mathcal{X}_{\text{train}}$, and Δx is a representation of the difference between x_{test} and x' . We then (ii) provide conditions under which $h(\Delta x, x')$ makes accurate predictions for unseen combinations of $(\Delta x, x')$ based on a theoretically justified bilinear modeling approach: $h(\Delta x, x') \rightarrow f(\Delta x)^\top g(x')$, where f and g map their inputs into vector spaces of the same dimension. Finally, (iii) empirical results demonstrate the generality of extrapolation of our algorithm on a wide variety of tasks: (a) regression for analytical functions and high-dimensional point cloud data; (b) sequential decision-making tasks such as goal-reaching for a simulated robot.

2 SETUP

Notation. Given a space of inputs \mathcal{X} and targets \mathcal{Y} , we aim to learn a predictor $h_\theta : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{Y})$ ¹ parameterized by θ , which best fits a ground truth function $h_* : \mathcal{X} \rightarrow \mathcal{Y}$. Given some non-negative loss function $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$ on the outputs (e.g., squared loss), and a distribution \mathcal{D} over \mathcal{X} , the *risk* is defined as

$$\mathcal{R}(h_\theta; \mathcal{D}) := \mathbb{E}_{x \sim \mathcal{D}} \mathbb{E}_{y \sim h_\theta(x)} \ell(y, h_*(x)). \quad (2.1)$$

Various training ($\mathcal{D}_{\text{train}}$) and test ($\mathcal{D}_{\text{test}}$) distributions yield different generalization settings:

In-Distribution Generalization. This setting assumes $\mathcal{D}_{\text{test}} = \mathcal{D}_{\text{train}}$. The challenge is to ensure that with N samples from $\mathcal{D}_{\text{train}}$, the expected risk $\mathcal{R}(h_\theta; \mathcal{D}_{\text{test}}) = \mathcal{R}(h_\theta; \mathcal{D}_{\text{train}})$ is small. This is a common paradigm in both empirical supervised learning (e.g., Simonyan & Zisserman (2014)) and in standard statistical learning theory (e.g., Vapnik (2006)).

Out-of-Distribution (OOD). This is more challenging and requires accurate predictions when $\mathcal{D}_{\text{train}} \neq \mathcal{D}_{\text{test}}$. When the ratio between the density function of $\mathcal{D}_{\text{test}}$ to that of $\mathcal{D}_{\text{train}}$ is bounded,

¹Throughout, we let $\mathcal{P}(\mathcal{Y})$ denote the set of distributions supported on \mathcal{Y} .

rigorous OOD extrapolation guarantees exist and are detailed in Appendix B.1. Such a situation arises when $\mathcal{D}_{\text{test}}$ shares support with $\mathcal{D}_{\text{train}}$ but is differently distributed as depicted in Fig 2a.

Out-of-Support (OOS). There are innumerable forms of distribution shift in which density ratios are not bounded. The most extreme case is when the support of $\mathcal{D}_{\text{test}}$ is not contained in that of $\mathcal{D}_{\text{train}}$. I.e., when there exists some $\mathcal{X}' \subset \mathcal{X}$ such that $\mathbb{P}_{x \sim \mathcal{D}_{\text{test}}}[x \in \mathcal{X}'] > 0$, but $\mathbb{P}_{x \sim \mathcal{D}_{\text{train}}}[x \in \mathcal{X}'] = 0$ (see Fig 2b). We term the problem of achieving low risk on such a $\mathcal{D}_{\text{test}}$ as *OOS extrapolation*.

Out-of-Combination (OOC). This is a special case of OOS. Let $\mathcal{X} = \mathcal{X}_1 \times \mathcal{X}_2$ be the product of two spaces. Let $\mathcal{D}_{\text{train}, \mathcal{X}_1}, \mathcal{D}_{\text{train}, \mathcal{X}_2}$ denote the marginal distributions of $x_1 \in \mathcal{X}_1, x_2 \in \mathcal{X}_2$ under $\mathcal{D}_{\text{train}}$, and $\mathcal{D}_{\text{test}, \mathcal{X}_1}, \mathcal{D}_{\text{test}, \mathcal{X}_2}$ under $\mathcal{D}_{\text{test}}$. In OOC learning, $\mathcal{D}_{\text{test}, \mathcal{X}_1}, \mathcal{D}_{\text{test}, \mathcal{X}_2}$ are in the support of $\mathcal{D}_{\text{train}, \mathcal{X}_1}, \mathcal{D}_{\text{train}, \mathcal{X}_2}$, but the joint distributions $\mathcal{D}_{\text{test}}$ need not be in the support of $\mathcal{D}_{\text{train}}$.

Transduction. In classical transduction (Gammerman et al., 1998), given data points $\{x_i\}_{i=1}^l$ and their labels $\{y_i\}_{i=1}^l$, the objective is making predictions for points $\{x_i\}_{i=l+1}^{l+k}$. In this paper we refer to predictors of x that are functions of the labeled data points $\{x_i\}_{i=1}^l$ as *transductive predictors*.




Figure 2: Illustration of different learning settings. (a) in-support out-of-distribution (OOD) learning; (b) general out-of-support (OOS) learning in 1-D (on the left) and 2-D (on the right); (c) out-of-combination (OOC) learning in 2-D.

3 CONDITIONS AND ALGORITHM FOR OUT-OF-SUPPORT GENERALIZATION

While general OOS extrapolation can be arbitrarily challenging, we show that there exist some concrete conditions under which OOC extrapolation is feasible. Under these conditions, an OOS prediction problem can be converted into an OOC prediction problem.




Figure 3: Illustration of converting OOS to OOC. (Left) Consider training points $x_1, x_2, x_3 \in \mathcal{X}_{\text{train}}$ and OOS test point x_{test} . During training, we predict $h_{\theta}(x_2)$ by transducing x_3 to $\bar{h}_{\theta}(\Delta_{23}, x_3)$, where $\Delta_{23} = x_2 - x_3$. Similarly, at test time, we predict $h_{\theta}(x_{\text{test}})$ by transducing training point x_1 , via $\bar{h}_{\theta}(\Delta_{x_{\text{test}}}, x_1)$, where $\Delta_{x_{\text{test}}} = x_{\text{test}} - x_1$. In this example note that $\Delta_{23} = \Delta_{x_{\text{test}}}$. (Right) This conversion yields an OOC generalization problem in space $\Delta \mathcal{X} \times \mathcal{X}_{\text{train}}$: marginal distributions $\Delta \mathcal{X}$ and $\mathcal{X}_{\text{train}}$ are covered by the training distribution, but their combination is not.

3.1 TRANSDUCTIVE PREDICTORS: CONVERTING OOS TO OOC

We require that input space \mathcal{X} has group structure, i.e., addition and subtraction operators $x+x'$, $x-x'$ are well-defined for $x, x' \in \mathcal{X}$. Let $\Delta\mathcal{X} := \{x-x' : x, x' \in \mathcal{X}\}$. We propose a *transductive reparameterization* $h_\theta : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{Y})$ with a *deterministic* function $\bar{h}_\theta : \Delta\mathcal{X} \times \mathcal{X} \rightarrow \mathcal{Y}$ as

$$h_\theta(x) := \bar{h}_\theta(x-x', x'), \quad (3.1)$$

where x' is referred to as an *anchor* point for a *query* point x .

Under this reparameterization, the **training** distribution can be rewritten as a *joint* distribution of $\Delta x = x-x'$ and x' as follows:

$$\mathbb{P}_{\bar{\mathcal{D}}_{\text{train}}}[(\Delta x, x') \in \cdot] := \Pr[(\Delta x, x') \in \cdot \mid x \sim \mathcal{D}_{\text{train}}, x' \sim \mathcal{D}_{\text{train}}, \Delta x = x-x']. \quad (3.2)$$

This is just representing the prediction for querying every point from the training distribution in terms of its relationship to other anchor points in the training distribution.

At **test** time, we are presented with query point $x \sim \mathcal{D}_{\text{test}}$ that may be from an OOS distribution. To make a prediction on this OOS x , we observe that with a careful selection of an anchor point x' from $\mathcal{D}_{\text{train}}$, our reparameterization may be able to convert this OOS problem into a more manageable OOC one, since representing the test point x in terms of its difference from training points can still be an “in-support” problem. For a radius parameter $\rho > 0$, define the distribution of chosen anchor points $\mathcal{D}_{\text{trns}}(x)$ (referred to as a *transducing* distribution) as

$$\mathbb{P}_{\mathcal{D}_{\text{trns}}(x)}[x' \in \cdot] = \Pr[x' \in \cdot \mid x' \sim \mathcal{D}_{\text{train}}, \inf_{\Delta x \in \Delta\mathcal{X}_{\text{train}}} \|(x-x') - \Delta x\| \leq \rho], \quad (3.3)$$

where $\mathcal{X}_{\text{train}}$ denotes the set of x in our training set, and we denote $\Delta\mathcal{X}_{\text{train}} := \{x_1-x_2 : x_1, x_2 \in \mathcal{X}_{\text{train}}\}$. Intuitively, our choice of $\mathcal{D}_{\text{trns}}(x)$ selects anchor points x' to transduce from the training distribution, subject to the resulting differences $(x-x')$ being close to a “seen” $\Delta x \in \Delta\mathcal{X}_{\text{train}}$. In doing so, both the anchor point x' and the difference Δx have been seen individually at training time, albeit not in combination. This allows us to express the prediction for a OOS query point in terms of an in-support anchor point x' and an in-support difference Δx (but *not jointly* in-support). This choice of anchor points induces a joint *test* distribution of $\Delta x = x-x'$ and x' :

$$\mathbb{P}_{\bar{\mathcal{D}}_{\text{test}}}[(\Delta x, x') \in \cdot] := \Pr[(\Delta x, x') \in \cdot \mid x \sim \mathcal{D}_{\text{test}}, x' \sim \mathcal{D}_{\text{trns}}(x), \Delta x = x-x']. \quad (3.4)$$

As seen from Fig 3, the marginals of Δx and x' under $\bar{\mathcal{D}}_{\text{test}}$, are individually in the support of those under $\bar{\mathcal{D}}_{\text{train}}$. Still, as Fig 3 reveals, since x_{test} is out-of-support, the joint distribution of $\bar{\mathcal{D}}_{\text{test}}$ is not covered by that of $\bar{\mathcal{D}}_{\text{train}}$ (i.e., the combination of x_1 and x_{test} have not been seen together before); precisely the OOC regime. Moreover, if one tried to transduce *all* $x' \sim \mathcal{D}_{\text{train}}$ to $x \sim \mathcal{D}_{\text{test}}$ at test time (e.g., transduce point x_3 to x_{test} in the figure) then we would lose coverage of the Δx -marginal. By applying transduction to keep both the marginal x' and Δx in-support, we are ensuring that we can convert difficult OOS problems into (potentially) more manageable OOC ones.

3.2 BILINEAR REPRESENTATIONS FOR OOC LEARNING

Without additional assumptions, OOC extrapolation may be just as challenging as OOS. However, with certain low-rank structure it can be feasible (Shah et al., 2020; Agarwal et al., 2021b; Athey et al., 2021). Following Agarwal et al. (2021a), we recognize that this low-rank property can be leveraged *implicitly* for our reparameterized OOC problem even in the continuous case (where $x, \Delta x$ do not explicitly form a finite dimensional matrix), using a bilinear representation of the *transductive* predictor in Eq. (3.1), $\bar{h}_\theta(\Delta x, x') = \langle f_\theta(\Delta x), g_\theta(x') \rangle$. Here f_θ, g_θ map their respective inputs into a vector space of the same dimension, \mathbb{R}^p .² If the output space is K -dimensional, then we independently model the prediction for each dimension using a set of K bilinear embeddings:

$$\bar{h}_\theta(\Delta x, x') = (\bar{h}_{\theta,1}(\Delta x, x'), \dots, \bar{h}_{\theta,K}(\Delta x, x')); \quad \bar{h}_{\theta,k}(\Delta x, x') = \langle f_{\theta,k}(\Delta x), g_{\theta,k}(x') \rangle. \quad (3.5)$$

While $\bar{h}_{\theta,k}$ are bilinear in embeddings $f_{\theta,k}, g_{\theta,k}$, the embeddings themselves may be parameterized by general function approximators. The effective “rank” of the transductive predictor is controlled by the dimension of the continuous embeddings $f_{\theta,k}(\Delta x), g_{\theta,k}(x')$.

²In our implementation, we take $\theta = (\theta_f, \theta_g)$, with separate parameters for each embedding.

We now introduce three conditions under which extrapolation is possible. As a preliminary, we write $\mu_1 \ll_{\kappa} \mu_2$ for two (unnecessarily normalized) measures μ_1, μ_2 to denote $\mu_1(A) \leq \kappa \mu_2(A)$ for all events A . This notion of coverage has been studied in the reinforcement learning community (termed as ‘‘concentrability’’ (Munos & Szepesvári, 2008)), and implies that the support of μ_1 is contained in the support of μ_2 . The first condition is the following notion of combinatorial coverage.

Assumption 3.1 (Bounded combinatorial density ratio). We assume that $\bar{\mathcal{D}}_{\text{test}}$ has κ -bounded combinatorial density ratio with respect to $\bar{\mathcal{D}}_{\text{train}}$, written as $\bar{\mathcal{D}}_{\text{test}} \ll_{\kappa, \text{comb}} \bar{\mathcal{D}}_{\text{train}}$. This means that there exist distributions $\mathcal{D}_{\Delta\mathcal{X}, i}$ and $\mathcal{D}_{\mathcal{X}, j}$, $i, j \in \{1, 2\}$, over $\Delta\mathcal{X}$ and \mathcal{X} , respectively, such that $\mathcal{D}_{i \otimes j} := \mathcal{D}_{\Delta\mathcal{X}, i} \otimes \mathcal{D}_{\mathcal{X}, j}$ satisfy

$$\sum_{(i,j) \neq (2,2)} \mathcal{D}_{i \otimes j} \ll_{\kappa} \bar{\mathcal{D}}_{\text{train}}, \quad \text{and} \quad \bar{\mathcal{D}}_{\text{test}} \ll_{\kappa} \sum_{i,j=1,2} \mathcal{D}_{i \otimes j}.$$

Assumption 3.1 can be interpreted as a generalization of a certain coverage condition in matrix factorization. The following distribution is elaborated upon in Appendix B.5 and Fig 9b therein. Intuitively, $\mathcal{D}_{1 \otimes 1}$ corresponds to the top-left block of a matrix, where all rows and columns are in-support. $\mathcal{D}_{1 \otimes 2}$ and $\mathcal{D}_{2 \otimes 1}$ correspond to the off-diagonals, and $\mathcal{D}_{2 \otimes 2}$ corresponds to the bottom-right block. The condition requires that $\bar{\mathcal{D}}_{\text{train}}$ covers the entire top-right block and off-diagonals, but need not cover the bottom-right block $\mathcal{D}_{2 \otimes 2}$. $\bar{\mathcal{D}}_{\text{test}}$, on the other hand, can contain samples from this last block. For problems with appropriate low-rank structure, samples from the top-left block, and samples from the off-diagonal blocks, uniquely determine the bottom-right block.

Recall that p denotes the dimension of the parametrized embeddings in Eq. (3.5). Our next assumptions are that the ground-truth embedding is low-rank (we refer to as ‘‘bilinearly transducible’’), and that under $\mathcal{D}_{1 \otimes 1}$, its component factors do not lie in a strict subspace of \mathbb{R}^p , i.e., not degenerate.

Assumption 3.2 (Bilinearly transducible). For each component $k \in [K]$, there exists $f_{*,k} : \Delta\mathcal{X} \rightarrow \mathbb{R}^p$ and $g_{*,k} : \mathcal{X} \rightarrow \mathbb{R}^p$ such that for all $x \in \mathcal{X}$, the following holds with probability 1 over $x' \sim \mathcal{D}_{\text{trns}}(x)$: $h_{*,k}(x) = \bar{h}_{*,k}(\Delta x, x') := \langle f_{*,k}(\Delta x), g_{*,k}(x') \rangle$, where $\Delta x = x - x'$. Further, the ground-truth predictions are bounded: $\max_k \sup_{\Delta x, x'} |\bar{h}_{*,k}(\Delta x, x')| \leq M$.

Assumption 3.3. For all components $k \in [K]$, the following lower bound holds for some $\sigma^2 > 0$:

$$\min\{\sigma_p(\mathbb{E}_{\mathcal{D}_{1 \otimes 1}}[f_{*,k} f_{*,k}^{\top}]), \sigma_p(\mathbb{E}_{\mathcal{D}_{1 \otimes 1}}[g_{*,k} g_{*,k}^{\top}])\} \geq \sigma^2. \quad (3.6)$$

The following is proven in Appendix B; due to limited space, we defer further discussion of the bound to Appendix B.4.1.

Theorem 1. Suppose Assumptions 3.1 to 3.3 all hold, and that the loss $\ell(\cdot, \cdot)$ is the square loss. Then, if the training risk satisfies $\mathcal{R}(h_{\theta}; \mathcal{D}_{\text{train}}) \leq \frac{\sigma^2}{4\kappa}$, then the test risk is bounded as

$$\mathcal{R}(h_{\theta}; \mathcal{D}_{\text{test}}) \leq \mathcal{R}(h_{\theta}; \mathcal{D}_{\text{train}}) \cdot \kappa^2 \left(1 + 64 \frac{M^4}{\sigma^4}\right) = \mathcal{R}(h_{\theta}; \mathcal{D}_{\text{train}}) \cdot \text{poly}(\kappa, \frac{M}{\sigma}).$$

3.3 OUR PROPOSED ALGORITHM: BILINEAR TRANSDUCTION

Our basic proposal for OOS extrapolation is **bilinear transduction**, depicted in Algorithm 1: at training time, a predictor \bar{h}_{θ} is trained to make predictions for training points x_i drawn from the training set $\mathcal{X}_{\text{train}}$ based on their similarity with other points $x_j \in \mathcal{X}_{\text{train}}$: $\bar{h}_{\theta}(x_i - x_j, x_j)$. The training pairs x_i, x_j are sampled uniformly from $\mathcal{X}_{\text{train}}$. At test time, for an OOS point x_{test} , we first select an anchor point x_i from the training set which has similarity with the test point $x_{\text{test}} - x_i$ that is within some radius ρ of the training similarities $\Delta\mathcal{X}_{\text{train}}$. We then predict the value for x_{test} based on the anchor point x_i and the similarity of the test and anchor points: $\bar{h}_{\theta}(x_{\text{test}} - x_i, x_i)$.

For supervised learning, i.e., the regression setting, we compute differences directly between inputs $x_i, x_j \in \mathcal{X}$. For the goal-conditioned imitation learning setting, we compute difference between states $x_i, x_j \in \mathcal{X}$ sampled uniformly over demonstration trajectories. At test time, we select an anchor trajectory based on the goal, and transduce each anchor state in the anchor trajectory to predict a sequence of actions for a test goal. As explained in Appendix C.1, there may be scenarios where it is beneficial to more carefully choose which points to transduce. To this end, we outline a more sophisticated proposal, **weighted transduction**, that achieves this. Pseudo code for bilinear transduction is depicted in Algorithm 1, and for the weighted variant in Algorithm 2 in the Appendix.

Note that we assume access to a representation of \mathcal{X} for which the subtraction operator captures differences that occur between training points and between training and test points. In Appendix D we describe examples of such representations as defined in standard simulators or which we extract with dimensionality reduction techniques (PCA). Alternatively, one might potentially extract such representations via self-supervision (Kingma & Welling, 2013) and pre-trained models (Upchurch et al., 2017), or compare data points via another similarity function.

Algorithm 1 Bilinear Transduction

- 1: **Input:** Distance parameter ρ , training set $(x_1, y_1), \dots, (x_n, y_n)$.
 - 2: **Train:** Train θ on loss $\mathcal{L}(\theta) = \sum_{i=1}^n \sum_{j \neq i} \ell(\bar{h}_\theta(x_i - x_j, x_j), y_i)$
 - 3: **Test:** For each new x_{test} , let $\mathcal{I}(x_{\text{test}}) := \{i : \inf_{\Delta x \in \Delta \mathcal{X}_{\text{train}}} \|x_{\text{test}} - x_i - \Delta x\| \leq \rho\}$, and predict

$$y = \bar{h}_\theta(x_{\text{test}} - x_i, x_i), \text{ where } i \sim \text{Uniform}(\mathcal{I}(x_{\text{test}}))$$
-

4 EXPERIMENTS

We answer the following questions through an empirical evaluation: **(1)** Does reformulating OOS extrapolation as a combinatorial generalization problem allow for extrapolation in a variety of supervised and sequential decision-making problems? **(2)** How does the particular choice of training distribution and data-generating function affect performance of the proposed technique? **(3)** How important is the choice of the low-rank bilinear function class for generalization? **(4)** Does our method scale to high dimensional state and action spaces? We first analyze OOS extrapolation on analytical domains and then on more real world problem domains depicted in Fig 8.

4.1 ANALYZING OOS EXTRAPOLATION ON ANALYTICAL PROBLEMS

We compare our method on regression problems generated via 1-D analytical functions (described in Appendix D.1) against standard neural networks with multiple fully-connected layers trained and tested in ranges $[20, 40]$ and $[10, 20] \cup [40, 50]$ (with the exception of Fig 4c trained in $[-1, 1]$ and tested in $[-1.6, -1] \cup [1, 1.6]$). We use these domains to gain intuition about the following questions:

What types of problems satisfy the assumptions for extrapolation? While we outlined a set of conditions under which extrapolation is guaranteed, it is not apparent which problems satisfy these assumptions. To understand this better, we considered learning functions with different structure: a periodic function with mixed periods (Fig 4a), a sawtooth function (Fig 4b) and a polynomial function (Fig 4c). Standard deep networks (yellow) fit the training points well (blue), but fail to extrapolate to OOS inputs (orange). In comparison, our approach (pink) accurately extrapolates on periodic functions but is much less effective on polynomials. This is because the periodic functions have symmetries which induce low rank structure under the proposed reparameterization.

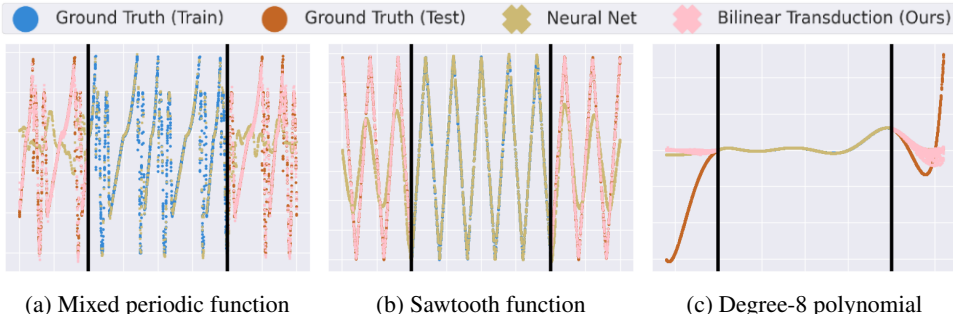


Figure 4: Bilinear transduction behavior on 1-D regression problems. Bilinear transduction performs well on functions with repeated structure, whereas they struggle on arbitrary polynomials. Standard neural nets fail to extrapolate in most settings, even when provided periodic activations (Tancik et al., 2020).

Going beyond known inductive biases. Given our method extrapolates to periodic functions in Fig 4 (which displays shift invariance), one might argue that a similar extrapolation can be achieved

by building in an inductive bias for periodicity / translation invariance. In Fig 6, we show that bilinear transduction in fact is able to extrapolate even in cases that the ground truth function is not simply translation invariant, but is translation equivariant, showing that bilinear transduction can capture equivariance. Moreover, bilinear transduction can in fact go beyond equivariance or invariance. To demonstrate broader generality of our method, we consider a piecewise periodic function that also grows in magnitude (Fig 7). This function is neither invariant nor equivariant to translation, as the group symmetries in this problem do not commute. The results in Fig 7 demonstrate that while the baselines fail to do so (including baselines that bake in equivariance (green)), bilinear transduction successfully extrapolates. The important thing for bilinear transduction to work is when comparing training instances, there is a simple (low-rank) relationship between how their labels are transformed. While it can capture invariance and equivariance, as Fig 7 shows, it is more general.

How does the relationship between the training distribution and test distribution affect extrapolation behavior? We try and understand how the range of test points that can be extrapolated to, depends on the training range. We show in Fig 5 that for a particular “width” of the training distribution (size of the training set), OOS extrapolation only extends for one “width” beyond the training range since the conditions for $\Delta\mathcal{X}$ being in-support are no longer valid beyond this point.




Figure 5: Performance of transductive predictors as test points go more and more OOS. Predictions are accurate for one “data-width” outside training data.

Figure 6: Prediction on function that displays affine equivariance. Bilinear transduction is able to capture equivariance without this being explicitly encoded.

Figure 7: Prediction on function that is neither invariant nor equivariant. Bilinear transduction is able to extrapolate while an equivariant predictor fails.

4.2 ANALYZING OOS EXTRAPOLATION ON LARGER SCALE DECISION-MAKING PROBLEMS

To establish that our method is useful for complex and real-world problem domains, we also vary the complexity (i.e., working with high-dimensional observation and action spaces) and the learning setting (regression and sequential decision making).




Figure 8: Evaluation domains at training (blue) and OOS (orange). (Left to Right:) grasp prediction for various object orientations and scales, table-top robotic manipulation for reaching and pushing objects to various targets, dexterous manipulation for relocating objects to various targets, slider control for striking a ball of various mass.

Baselines. To assess the effectiveness of our proposed scheme for extrapolation via transduction and bilinear embeddings, we compare with the following non-transductive baselines. *Linear Model:* linear function approximator to understand whether the issue is one of overparameterization, and whether linear models would solve the problem. *Neural Networks:* typical training of overparameterized neural network function approximator via standard empirical risk minimization. *Alternative Techniques with Neural Networks (DeepSets):* an alternative architecture for combining multiple inputs (DeepSets (Zaheer et al., 2017)), that are meant to be permutation invariant and encourage a degree of generaliza-

tion between different pairings of states and goals. Finally, we compare with a *Transductive Method without a Mechanism for Structured Extrapolation (Transduction)*: transduction with no special structure for understanding the impact of bilinear embeddings and the low-rank structure. This baseline uses reparameterization, but just parameterizes \bar{h}_θ as a standard neural network. We present additional comparison results introducing periodic activations (Tancik et al., 2020) in Table 4 in the Appendix.

OOS extrapolation in sequential decision making. Table 1 contains results for different types of extrapolation settings. More training and evaluation details are provided in Appendices D.1 and D.2.

- *Extrapolation to OOS Goals:* We considered two tasks from the Meta-World benchmark (Yu et al., 2020) where a simulated robotic agent needs to either reach or push a target object to a goal location (column 2 in Fig 8). Given a set of expert demonstrations reaching/pushing to goals in the blue box ($[0, 0.4] \times [0.7, 0.9] \times [0.05, 0.3]$ and $[0, 0.3] \times [0.5, 0.7] \times \{0.01\}$), we tested generalization to OOS goals in the orange box ($[-0.4, 0] \times [0.7, 0.9] \times [0.05, 0.3]$ and $[-0.3, 0] \times [0.5, 0.7] \times \{0.01\}$), using a simple extension of our method described in Appendix D.2 to perform transduction over trajectories rather than individual states. We quantify performance by measuring the distance between the conditioned and reached goal. Results in Table 1 show that on the easy task of reaching, training a linear or typical neural network-based predictor extrapolate as well as our method. However, for the more challenging task of pushing an object, our extrapolation is better by an order of magnitude than other baselines, showing the ability to generalize goals in a completely different direction.
- *Extrapolation with Large State and Action Space:* Next we tested our method on grasping and placing an object to OOS goal-locations in \mathbb{R}^3 with an anthropomorphic “Adroit” hand with a much larger action (\mathbb{R}^{30}) and state (\mathbb{R}^{39}) space (column 3 in Fig 8). Results confirm that bilinear transduction scales up to high dimensional state-action spaces as well and is naturally able to grasp the ball and move it to new target locations ($[-0.3, 0] \times [-0.3, 0] \times [0.15, 0.35]$) after trained on target locations in ($[0, 0.3] \times [0, 0.3] \times [0.15, 0.35]$). These results are better appreciated in the video attached with the supplementary material, but show the same trends as above, with bilinear transduction significantly outperforming standard inductive methods and non-bilinear architectures.
- *Extrapolation to OOS Dynamics:* Lastly, we consider problems involving extrapolation not just in location of the goals for goal-reaching problems, but across problems with varying dynamics. Specifically, we consider a slider task where the goal is to move a slider on a table to strike a ball such that it rolls to a fixed target position (column 4 in Fig 8). The mass of the ball varies across episodes and is provided as input to the policy. We train and test on a range of masses ($[60, 130]$ and $[5, 15]$). We find that bilinear transduction adjusts behavior and successfully extrapolates to new masses, showing the ability to extrapolate not just to goals, but also to varying dynamics.

Importantly, bilinear transduction is significantly less prone to variance than standard inductive or permutation-invariant architectures. This can be seen from a heatmap over various training architectures and training seeds as shown in Fig 12 and Table 5 in Appendix C.4. While other methods can sometimes show success, only bilinear transduction is able to consistently accomplish extrapolation.

OOS extrapolation in higher dimensional regression problems. To scale up the dimension of the input space, we consider the problem of predicting valid grasping points in \mathbb{R}^3 from point clouds of various objects (bottles, mugs, and teapots) with different orientations, positions, and scales (column 1 in Fig 8). At training and test, objects undergo z -axis orientation ($[0, 1.2\pi]$ and $[1.2\pi, 2\pi]$), translation ($[0, 0.5] \times [0, 0.5]$ and $[0.5, 0.7] \times [0.5, 0.7]$) and scaling ($[0.7, 1.3]$ and $[1.3, 1.6]$). In this domain, we represent entire point clouds by a low-dimensional representation of the point cloud obtained via PCA. We consider situations, where we train on individual objects where the training set consists of various rotations, translations, and scales, and standard bilinear transduction is applied. We also consider situations where the objects are not individually identified but instead, a single grasp point predictor is trained on the entire set of bottles, mugs, and teapots. We assume access to category labels at training time, but do not require this at test time. For a more in-depth discussion of this assumption and a discussion of how this can be learned via *weighted*-transduction, we refer readers to Appendix C.1. While training performance is comparable in all instances, we find that extrapolation behavior is significantly better for our method. This is true both for single object cases as well as scenarios where all objects are mixed together (Table 1). These experiments show that bilinear transduction can work on feature spaces of high-dimensional data such as point clouds. Note that while we only predict a single grasp point here, a single grasp point prediction can be easily generalized to predict multiple keypoints instead (Manuelli et al., 2019), enabling success on more challenging control domains. For further visualizations and details, please see Appendices C and D.

Table 1: Mean and standard deviation over prediction (regression) or final state (sequential decision making) error for OOS samples and over a hyperparameter search.

Task	Expert	Linear	Neural Net	DeepSets	Transduction	Ours
Mug		0.068 ± 0.013	0.075 ± 0.046		0.055 ± 0.043	0.013 ± 0.007
Bottle		0.026 ± 0.005	0.05 ± 0.051		0.027 ± 0.016	0.008 ± 0.004
Teapot		0.095 ± 0.02	0.101 ± 0.078		0.043 ± 0.02	0.022 ± 0.014
All		0.143 ± 0.116	0.118 ± 0.075		0.112 ± 0.08	0.018 ± 0.012
Reach	0.006 ± 0.008	0.007 ± 0.006	0.036 ± 0.054	0.19 ± 0.209	0.036 ± 0.048	0.007 ± 0.006
Push	0.012 ± 0.001	0.258 ± 0.063	0.258 ± 0.167	0.199 ± 0.114	0.159 ± 0.116	0.02 ± 0.017
Slider	0.105 ± 0.066	0.609 ± 0.07	0.469 ± 0.336	0.274 ± 0.262	0.495 ± 0.339	0.149 ± 0.113
Adroit	0.035 ± 0.015	0.337 ± 0.075	0.331 ± 0.203	0.521 ± 0.457	0.409 ± 0.32	0.147 ± 0.117

5 ABRIDGED RELATED WORK

Generalization from training data to test data of the same distribution has been studied extensively, both practically (Simonyan & Zisserman, 2014) and theoretically (Vapnik, 2006; Bousquet & Elisseeff, 2002; Bartlett & Mendelson, 2002). Our focus is on performance on distributions with examples which may not be covered by the training data, as described formally in Section 2. Here we provide a discussion of the most directly related approaches, and defer an extended discussion to Appendix A. Even more so than generalization, extrapolation necessitates leveraging structure in both the data and the learning algorithm. Along these lines, past work has focused on *structured neural networks*, which hardcode such symmetries as equivariance (Cohen & Welling, 2016; Simeonov et al., 2022), Euclidean symmetry (Smidt, 2021) and periodicity (Abu-Dakka et al., 2021; Parascandolo et al., 2016) into the learning process. Other directions have focused on general *architectures* which seem to exhibit combinatorial generalization, such as transformers (Vaswani et al., 2017), graph neural networks (Cappart et al., 2021) and bilinear models (Hong et al., 2021).

In this work, we focus on bilinear architectures with what we term a *transductive parametrization*. We demonstrate that this framework can in many cases learn the symmetries (e.g. equivariance and periodicity) that structured neural networks hardcode for, and achieve extrapolation in some regimes in which these latter methods cannot. A bilinear model with low inner dimension is equivalent to enforcing a low-rank constraint on one’s predictions. Low-rank models have commanded broad popularity for matrix completion (Mnih & Salakhutdinov, 2007; Mackey et al., 2015). And whereas early analysis focused on the *missing-at-random* setting (Candès & Recht, 2009; Recht, 2011; Candès & Plan, 2010) equivalent to classical in-distribution statistical learning, we adopt more recent perspectives on missing-not-at-random data, see e.g., (Shah et al., 2020; Agarwal et al., 2021b; Athey et al., 2021), which tackle the combinatorial-generalization setting described in Section 3.2.

In the classical *transduction* setting (Joachims, 2003; Gammelman et al., 1998; Cortes & Mohri, 2006), the goal is to make predictions on a known set of *unlabeled* test examples for which the features are known; this is a special case of semi-supervised learning. We instead operate in a standard supervised learning paradigm, where test labels *and* features are only revealed at test time. Still, we find it useful to adopt a “transductive parametrization” of our predictor, where instead of compressing our predictor into parameters alone, we express predictions for labels of one example as a function of other, labeled examples. An unabridged version of related work is in Appendix A.

6 DISCUSSION

Our work serves as an initial study of the circumstances under which problem structure can be both *discovered* and *exploited* for extrapolation, combining parametric and non-parametric approaches. The main limitations of this work are assumptions regarding access to a representation of \mathcal{X} and similarity measure for obtaining Δx . Furthermore, we are only guaranteed to extrapolate to regions of \mathcal{X} that admit Δx within the training distribution. A number of natural questions arise for further research. First, can we classify which set of real-world domains fits our assumptions, beyond the domains we have demonstrated? Second, can we *learn* a representation of \mathcal{X} in which differences Δx are meaningful for high dimensional domains? And lastly, are there more effective schemes for selecting anchor points? For instance, analogy-making for anchor point selection may reduce the complexity of \bar{h}_θ , guaranteeing low-rank structure.

7 ETHICS STATEMENT

Bias In this work on extrapolation via bilinear transduction, we can only aim to extrapolate to data that shifts from the training set in ways that exist within the training data. Therefore the performance of this method relies on the diversity within the training data, allowing some forms of extrapolation but not to ones that do not exist within the training distribution.

Dataset release We provide in detail in the Appendix the parameters and code bases and datasets we use to generate our data. In the future, we plan to release our code base and expert policy weights with which we collected expert data.

8 REPRODUCIBILITY STATEMENT

We describe our algorithms in Section 3.3 and the complete proof of our theoretical results and assumptions in Appendix B. Extensive implementation details regarding our algorithms, data, models, and optimization are provided in Appendix D. In addition, we plan in the future to release our code and data.

9 ACKNOWLEDGMENTS

We thank Anurag Ajay, Tao Chen, Zhang-Wei Hong, Jacob Huh, Leslie Kaelbling, Hannah Lawrence, Richard Li, Gabe Margolis, Devavrat Shah and Anthony Simeonov for the helpful discussions and feedback on the paper. We are grateful to MIT Supercloud and the Lincoln Laboratory Supercomputing Center for providing HPC resources. This research was also partly sponsored by the DARPA Machine Common Sense Program, MIT-IBM Watson AI Lab, the United States Air Force Research Laboratory and the United States Air Force Artificial Intelligence Accelerator and was accomplished under Cooperative Agreement Number FA8750-19-2-1000. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the United States Air Force or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes, notwithstanding any copyright notation herein.

REFERENCES

- Fares J Abu-Dakka, Matteo Saveriano, and Luka Peternel. Periodic dmp formulation for quaternion trajectories. In *2021 20th International Conference on Advanced Robotics (ICAR)*, pp. 658–663. IEEE, 2021.
- Anish Agarwal, Abdullah Alomar, Varkey Alumootil, Devavrat Shah, Dennis Shen, Zhi Xu, and Cindy Yang. Persim: Data-efficient offline reinforcement learning with heterogeneous agents via personalized simulators. *Advances in Neural Information Processing Systems*, 34:18564–18576, 2021a.
- Anish Agarwal, Munther Dahleh, Devavrat Shah, and Dennis Shen. Causal matrix completion. *arXiv preprint arXiv:2109.15154*, 2021b.
- Jacob Andreas. Measuring compositionality in representation learning. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL <https://openreview.net/forum?id=HJz05o0qK7>.
- Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Dan Klein. Neural module networks. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pp. 39–48. IEEE Computer Society, 2016. doi: 10.1109/CVPR.2016.12. URL <https://doi.org/10.1109/CVPR.2016.12>.
- Marcin Andrychowicz, Dwight Crow, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, Pieter Abbeel, and Wojciech Zaremba. Hindsight experience replay. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 5048–5058, 2017. URL <https://proceedings.neurips.cc/paper/2017/hash/453fadbd8a1a3af50a9df4df899537b5-Abstract.html>.
- Susan Athey, Mohsen Bayati, Nikolay Doudchenko, Guido Imbens, and Khashayar Khosravi. Matrix completion methods for causal panel data models. *Journal of the American Statistical Association*, 116(536):1716–1730, 2021.
- Andrei Barbu, David Mayo, Julian Alverio, William Luo, Christopher Wang, Dan Gutfreund, Josh Tenenbaum, and Boris Katz. Objectnet: A large-scale bias-controlled dataset for pushing the limits of object recognition models. *Advances in neural information processing systems*, 32, 2019.
- Peter L Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002.
- Peter W Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018.
- Gregory Benton, Marc Finzi, Pavel Izmailov, and Andrew G Wilson. Learning invariances in neural networks from training data. *Advances in neural information processing systems*, 33:17605–17616, 2020.
- Olivier Bousquet and André Elisseeff. Stability and generalization. *The Journal of Machine Learning Research*, 2:499–526, 2002.
- Emmanuel J Candes and Yaniv Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925–936, 2010.
- Emmanuel J Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717–772, 2009.
- Quentin Cappart, Didier Chételat, Elias Khalil, Andrea Lodi, Christopher Morris, and Petar Veličković. Combinatorial optimization and reasoning with graph neural networks. *arXiv preprint arXiv:2102.09544*, 2021.

- Rich Caruana. Multitask learning. *Mach. Learn.*, 28(1):41–75, 1997. doi: 10.1023/A:1007379606734. URL <https://doi.org/10.1023/A:1007379606734>.
- Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.
- Erhan Çinlar. *Probability and stochastics*, volume 261. Springer, 2011.
- Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International conference on machine learning*, pp. 2990–2999. PMLR, 2016.
- Corinna Cortes and Mehryar Mohri. On transductive regression. *Advances in neural information processing systems*, 19, 2006.
- Michael Crawshaw. Multi-task learning with deep neural networks: A survey. *CoRR*, abs/2009.09796, 2020. URL <https://arxiv.org/abs/2009.09796>.
- Pim de Haan, Dinesh Jayaraman, and Sergey Levine. Causal confusion in imitation learning. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 11693–11704, 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/947018640bf36a2bb609d3557a285329-Abstract.html>.
- Nima Dehmamy, Robin Walters, Yan Chen Liu, Dashun Wang, and Rose Yu. Automatic symmetry discovery with lie algebra convolutional network. *Advances in Neural Information Processing Systems*, 34:2503–2515, 2021.
- Congyue Deng, Or Litany, Yueqi Duan, Adrien Poulenard, Andrea Tagliasacchi, and Leonidas J Guibas. Vector neurons: A general framework for so(3)-equivariant networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 12200–12209, 2021a.
- Congyue Deng, Or Litany, Yueqi Duan, Adrien Poulenard, Andrea Tagliasacchi, and Leonidas J. Guibas. Vector neurons: A general framework for so(3)-equivariant networks. *CoRR*, abs/2104.12229, 2021b. URL <https://arxiv.org/abs/2104.12229>.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Benjamin L Edelman, Surbhi Goel, Sham Kakade, and Cyril Zhang. Inductive biases and variable creation in self-attention mechanisms. In *International Conference on Machine Learning*, pp. 5793–5831. PMLR, 2022.
- Alireza Fallah, Aryan Mokhtari, and Asuman E. Ozdaglar. On the convergence theory of gradient-based model-agnostic meta-learning algorithms. In Silvia Chiappa and Roberto Calandra (eds.), *The 23rd International Conference on Artificial Intelligence and Statistics, AISTATS 2020, 26-28 August 2020, Online [Palermo, Sicily, Italy]*, volume 108 of *Proceedings of Machine Learning Research*, pp. 1082–1092. PMLR, 2020. URL <http://proceedings.mlr.press/v108/fallah20a.html>.
- Angelos Filos, Panagiotis Tigas, Rowan McAllister, Nicholas Rhinehart, Sergey Levine, and Yarın Gal. Can autonomous vehicles identify, recover from, and adapt to distribution shifts? In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pp. 3145–3153. PMLR, 2020. URL <http://proceedings.mlr.press/v119/filos20a.html>.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. *CoRR*, abs/1703.03400, 2017a. URL <http://arxiv.org/abs/1703.03400>.

- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pp. 1126–1135. PMLR, 2017b.
- Alex Gammerman, Volodya Vovk, and Vladimir Vapnik. Learning by transduction. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, pp. 148–155, 1998.
- Dibya Ghosh, Jad Rahme, Aviral Kumar, Amy Zhang, Ryan P. Adams, and Sergey Levine. Why generalization in RL is difficult: Epistemic pomdps and implicit partial observability. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pp. 25502–25515, 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/d5fff135377d39f1de7372c95c74dd962-Abstract.html>.
- Alex Gittens and Michael Mahoney. Revisiting the nystrom method for improved large-scale machine learning. In *International Conference on Machine Learning*, pp. 567–575. PMLR, 2013.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In Jennifer G. Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 1856–1865. PMLR, 2018. URL <http://proceedings.mlr.press/v80/haarnoja18b.html>.
- Philippe Hansen-Estruch, Amy Zhang, Ashvin Nair, Patrick Yin, and Sergey Levine. Bisimulation makes analogies in goal-conditioned reinforcement learning. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato (eds.), *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pp. 8407–8426. PMLR, 2022. URL <https://proceedings.mlr.press/v162/hansen-estruch22a.html>.
- Zhang-Wei Hong, Ge Yang, and Pulkit Agrawal. Bi-linear value networks for multi-goal reinforcement learning. In *International Conference on Learning Representations*, 2021.
- Nicholas Ichien, Qing Liu, Shuhao Fu, Keith J. Holyoak, Alan L. Yuille, and Hongjing Lu. Visual analogy: Deep learning versus compositional models. *CoRR*, abs/2105.07065, 2021. URL <https://arxiv.org/abs/2105.07065>.
- Stefanie Jegelka. Theory of graph neural networks: Representation and learning. *arXiv preprint arXiv:2204.07697*, 2022.
- Thorsten Joachims. Transductive learning via spectral graph partitioning. In *Proceedings of the 20th international conference on machine learning (ICML-03)*, pp. 290–297, 2003.
- Ryan Julian, Benjamin Swanson, Gaurav S. Sukhatme, Sergey Levine, Chelsea Finn, and Karol Hausman. Efficient adaptation for end-to-end vision-based robotic manipulation. *CoRR*, abs/2004.10190, 2020. URL <https://arxiv.org/abs/2004.10190>.
- Leslie Pack Kaelbling. Learning to achieve goals. In Ruzena Bajcsy (ed.), *Proceedings of the 13th International Joint Conference on Artificial Intelligence. Chambéry, France, August 28 - September 3, 1993*, pp. 1094–1099. Morgan Kaufmann, 1993.
- Dmitry Kalashnikov, Jacob Varley, Yevgen Chebotar, Benjamin Swanson, Rico Jonschkowski, Chelsea Finn, Sergey Levine, and Karol Hausman. Mt-opt: Continuous multi-task robotic reinforcement learning at scale. *CoRR*, abs/2104.08212, 2021. URL <https://arxiv.org/abs/2104.08212>.
- Osman Semih Kayhan and Jan C van Gemert. On translation invariance in cnns: Convolutional layers can exploit absolute spatial location. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14274–14285, 2020.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Robert Kirk, Amy Zhang, Edward Grefenstette, and Tim Rocktäschel. A survey of generalisation in deep reinforcement learning. *CoRR*, abs/2111.09794, 2021. URL <https://arxiv.org/abs/2111.09794>.
- Nikita Kitaev, Łukasz Kaiser, and Anselm Levskaya. Reformer: The efficient transformer. *arXiv preprint arXiv:2001.04451*, 2020.
- Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, Tony Lee, Etienne David, Ian Stavness, Wei Guo, Berton Earnshaw, Imran S. Haque, Sara M. Beery, Jure Leskovec, Anshul Kundaje, Emma Pierson, Sergey Levine, Chelsea Finn, and Percy Liang. WILDS: A benchmark of in-the-wild distribution shifts. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pp. 5637–5664. PMLR, 2021. URL <http://proceedings.mlr.press/v139/koh21a.html>.
- Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, et al. Solving quantitative reasoning problems with language models. *arXiv preprint arXiv:2206.14858*, 2022.
- Lester W Mackey, Ameet Talwalkar, and Michael I Jordan. Distributed matrix completion and robust factorization. *J. Mach. Learn. Res.*, 16(1):913–960, 2015.
- Lucas Manuelli, Wei Gao, Peter R. Florence, and Russ Tedrake. KPAM: keypoint affordances for category-level robotic manipulation. In Tamim Asfour, Eiichi Yoshida, Jaeheung Park, Henrik Christensen, and Oussama Khatib (eds.), *Robotics Research - The 19th International Symposium ISRR 2019, Hanoi, Vietnam, October 6-10, 2019*, volume 20 of *Springer Proceedings in Advanced Robotics*, pp. 132–157. Springer, 2019. doi: 10.1007/978-3-030-95459-8_9. URL https://doi.org/10.1007/978-3-030-95459-8_9.
- Sara Meftah, Nasredine Semmar, Mohamed-Ayoub Tahiri, Youssef Tamaazousti, Hassane Essafi, and Fatiha Sadat. Multi-task supervised pretraining for neural domain adaptation. In *Proceedings of the Eighth International Workshop on Natural Language Processing for Social Media*, pp. 61–71, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.socialnlp-1.8. URL <https://aclanthology.org/2020.socialnlp-1.8>.
- Lingsheng Meng and Bing Zheng. The optimal perturbation bounds of the moore–penrose inverse under the frobenius norm. *Linear algebra and its applications*, 432(4):956–963, 2010.
- Melanie Mitchell. Abstraction and analogy-making in artificial intelligence. *CoRR*, abs/2102.10717, 2021. URL <https://arxiv.org/abs/2102.10717>.
- Andriy Mnih and Russ R Salakhutdinov. Probabilistic matrix factorization. *Advances in neural information processing systems*, 20, 2007.
- Rémi Munos and Csaba Szepesvári. Finite-time bounds for fitted value iteration. *Journal of Machine Learning Research*, 9(5), 2008.
- Ashvin Nair, Vitchyr Pong, Murtaza Dalal, Shikhar Bahl, Steven Lin, and Sergey Levine. Visual reinforcement learning with imagined goals. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 9209–9220, 2018. URL <https://proceedings.neurips.cc/paper/2018/hash/7ec69dd44416c46745f6edd947b470cd-Abstract.html>.
- Giambattista Parascandolo, Heikki Huttunen, and Tuomas Virtanen. Taming the waves: sine as activation function in deep neural networks. 2016.
- Joaquin Quinonero-Candela, Masashi Sugiyama, Anton Schwaighofer, and Neil D Lawrence. *Dataset shift in machine learning*. Mit Press, 2008.

- Hamed Rahimian and Sanjay Mehrotra. Distributionally robust optimization: A review, 2019. URL <https://arxiv.org/abs/1908.05659>.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pp. 8821–8831. PMLR, 2021.
- Benjamin Recht. A simpler approach to matrix completion. *Journal of Machine Learning Research*, 12(12), 2011.
- Scott E. Reed, Yi Zhang, Yuting Zhang, and Honglak Lee. Deep visual analogy-making. In Corinna Cortes, Neil D. Lawrence, Daniel D. Lee, Masashi Sugiyama, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pp. 1252–1260, 2015. URL <https://proceedings.neurips.cc/paper/2015/hash/e07413354875be01a996dc560274708e-Abstract.html>.
- Clemens Rosenbaum, Tim Klinger, and Matthew Riemer. Routing networks: Adaptive selection of non-linear functions for multi-task learning. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. URL <https://openreview.net/forum?id=ry8dvM-R->.
- Adam Santoro, Sergey Bartunov, Matthew M. Botvinick, Daan Wierstra, and Timothy P. Lillicrap. One-shot learning with memory-augmented neural networks. *CoRR*, abs/1605.06065, 2016. URL <http://arxiv.org/abs/1605.06065>.
- John Schulman, Sergey Levine, Pieter Abbeel, Michael I. Jordan, and Philipp Moritz. Trust region policy optimization. In Francis R. Bach and David M. Blei (eds.), *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pp. 1889–1897. JMLR.org, 2015. URL <http://proceedings.mlr.press/v37/schulman15.html>.
- Devavrat Shah, Dogyoon Song, Zhi Xu, and Yuzhe Yang. Sample efficient reinforcement learning via low-rank matrix estimation. *Advances in Neural Information Processing Systems*, 33:12092–12103, 2020.
- Anthony Simeonov, Yilun Du, Andrea Tagliasacchi, Joshua B Tenenbaum, Alberto Rodriguez, Pulkit Agrawal, and Vincent Sitzmann. Neural descriptor fields: Se (3)-equivariant object representations for manipulation. In *2022 International Conference on Robotics and Automation (ICRA)*, pp. 6394–6400. IEEE, 2022.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Aman Sinha, Hongseok Namkoong, and John C. Duchi. Certifying some distributional robustness with principled adversarial training. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. URL <https://openreview.net/forum?id=Hk6kPgZA->.
- Tess E Smidt. Euclidean symmetry and equivariance in machine learning. *Trends in Chemistry*, 3(2):82–85, 2021.
- Gilbert W Stewart. On the perturbation of pseudo-inverses, projections and linear least squares problems. *SIAM review*, 19(4):634–662, 1977.
- Yixuan Su, Lei Shu, Elman Mansimov, Arshit Gupta, Deng Cai, Yi-An Lai, and Yi Zhang. Multi-task pre-training for plug-and-play task-oriented dialogue system. In Smaranda Muresan, Preslav Nakov, and Aline Villavicencio (eds.), *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pp. 4661–4676. Association for Computational Linguistics, 2022. doi: 10.18653/v1/2022.acl-long.319. URL <https://doi.org/10.18653/v1/2022.acl-long.319>.

- Matthew Tancik, Pratul P. Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T. Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *CoRR*, abs/2006.10739, 2020. URL <https://arxiv.org/abs/2006.10739>.
- Nathaniel Thomas, Tess Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley. Tensor field networks: Rotation-and translation-equivariant neural networks for 3d point clouds. *arXiv preprint arXiv:1802.08219*, 2018.
- Sebastian Thrun and Lorien Y. Pratt (eds.). *Learning to Learn*. Springer, 1998. ISBN 978-1-4613-7527-2. doi: 10.1007/978-1-4615-5529-2. URL <https://doi.org/10.1007/978-1-4615-5529-2>.
- Paul Upchurch, Jacob Gardner, Geoff Pleiss, Robert Pless, Noah Snavely, Kavita Bala, and Kilian Weinberger. Deep feature interpolation for image content changes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7064–7073, 2017.
- Vladimir Vapnik. *Estimation of dependences based on empirical data*. Springer Science & Business Media, 2006.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Oriol Vinyals, Charles Blundell, Timothy P. Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. Matching networks for one shot learning. *CoRR*, abs/1606.04080, 2016. URL <http://arxiv.org/abs/1606.04080>.
- Dian Wang, Robin Walters, and Robert Platt. $\text{SO}(2)$ -equivariant reinforcement learning. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022. URL https://openreview.net/forum?id=7F9cOhdvfk_.
- Mingzhang Yin, George Tucker, Mingyuan Zhou, Sergey Levine, and Chelsea Finn. Meta-learning without memorization. *CoRR*, abs/1912.03820, 2019. URL <http://arxiv.org/abs/1912.03820>.
- Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Karol Hausman, Chelsea Finn, and Sergey Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *Conference on Robot Learning*, pp. 1094–1100. PMLR, 2020.
- Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R Salakhutdinov, and Alexander J Smola. Deep sets. *Advances in neural information processing systems*, 30, 2017.
- Jian Zhang, Yuanqing Zhang, Huan Fu, Xiaowei Zhou, Bowen Cai, Jinchu Huang, Rongfei Jia, Binqiang Zhao, and Xing Tang. Ray priors through reprojection: Improving neural radiance fields for novel view extrapolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18376–18386, 2022a.
- Yi Zhang, Arturs Backurs, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, and Tal Wagner. Unveiling transformers with lego: a synthetic reasoning task. *arXiv preprint arXiv:2206.04301*, 2022b.
- Wenshuai Zhao, Jorge Peña Queraltá, and Tomi Westerlund. Sim-to-real transfer in deep reinforcement learning for robotics: a survey. In *2020 IEEE symposium series on computational intelligence (SSCI)*, pp. 737–744. IEEE, 2020.
- Allan Zhou, Vikash Kumar, Chelsea Finn, and Aravind Rajeswaran. Policy architectures for compositional generalization in control. *arXiv preprint arXiv:2203.05960*, 2022.

CONTENTS

1	Introduction	1
2	Setup	2
3	Conditions and Algorithm for Out-of-Support Generalization	3
3.1	Transductive Predictors: Converting OOS to OOC	4
3.2	Bilinear representations for OOC learning	4
3.3	Our Proposed Algorithm: Bilinear Transduction	5
4	Experiments	6
4.1	Analyzing OOS extrapolation on analytical problems	6
4.2	Analyzing OOS extrapolation on larger scale decision-making problems	7
5	Abridged Related Work	9
6	Discussion	9
7	Ethics Statement	10
8	Reproducibility Statement	10
9	Acknowledgments	10
A	Unabridged Related Work	19
B	Theoretical Results	21
B.1	Generalization under bounded density ratio	21
B.2	Extrapolation for Matrix Completion	22
B.3	General Analysis for Combinatorial Extrapolation	22
B.3.1	Proof of Theorem 2	23
B.4	Extrapolation for Transduction	25
B.4.1	Further Remarks on Theorem 1	26
B.5	Connections to matrix completion	26
C	Additional Results	27
C.1	Learning which points to transduce: Weighted Transduction	27
C.2	Additional Results on a Simple 2-D Example	27
C.3	Additional Results on Supervised Grasp prediction	28
C.4	Additional Results on Imitation Learning	29
C.5	Choice of Architecture	30
D	Implementation Details	31

D.1	Train/Test distributions and Analytic Function Details	31
D.2	Training details	32
D.2.1	Details of Bilinear Transduction Implementation	32
D.2.2	Details of Training Distributions	33
D.2.3	Details of Model Architectures	34
D.2.4	Details of Model Optimization	34

A UNABRIDGED RELATED WORK

Here we discuss various prior approaches to extrapolation in greater depth, focusing on areas not addressed in our abridged discussion in the main text.

Approaches which encode explicit structure. One popular approach to designing networks that extrapolate to novel data has been the *equivariant neural networks*, first proposed by [Cohen & Welling \(2016\)](#). The key idea is that, if it is known there is a group G which acts on both the input and target domains of the predictor, and if it is understood that the true predictor must satisfy the *equivariance property*

$$h_*(g \cdot x) = g \cdot h_*(x)$$

for all $g \in G$ (here $g \cdot ()$ denotes group action), then one can explicitly encode for predictors h_θ satisfying the same identity. Similarly, one can encode for invariances, where $h_*(g \cdot x) = h_*(x)$ for all $g \in G$.

[Deng et al. \(2021a\)](#) extended the original equivariance framework to accommodate situations where the group G corresponds to rotations ($SO(3)$), and [Simeonov et al. \(2022\)](#) proposed neural descriptor fields which handle rigid-body transformations $SE(3)$. For a broader review on how other notions of symmetry can be encoded into machine learning settings, consult [Smidt \(2021\)](#), and [Abu-Dakka et al. \(2021\)](#); [Parascandolo et al. \(2016\)](#) for how periodic structure can be explicitly built in. We remark that in many of these approaches, the group/symmetry must be known ahead of time. While attempt to learn global ([Benton et al., 2020](#)) or local ([Dehmamy et al., 2021](#)) equivari-ances/invariances, there are numerous forms of structure that can be represented as a group symme-try, and for which these methods do not apply. As we show in our experiments, there are a number of group structures which are *not* captured by equivariance (such as Fig 7), which bilinear transduction captures.

Architectures favorable to extrapolation. There are numerous learning architectures which are purported to be favorable to extrapolation to novel-domains. For example, graph neural networks (GNNs) have been used to facilitate reasoning behavior in combinatorial environments ([Battaglia et al., 2018](#); [Cappart et al., 2021](#)), and the implicit biases of GNNs have received much theoretical study in recent years (see [Jegelka \(2022\)](#) and the references therein). Another popular domain for extrapolations has been sequence-to-sequence modeling, especially in the context of natural language processing. Here, the now-renowned Transformer architecture due to [Vaswani et al. \(2017\)](#), as well as its variants based on the same “attention mechanism” (e.g. [Kitaev et al. \(2020\)](#)) have become incredibly popular. Supervised by a sufficiently diverse set of tasks and staggering amount of data, these have shown broad population in various language understanding tasks ([Devlin et al., 2018](#)), text-to-image generation ([Ramesh et al., 2021](#)), and even quantitative reasoning ([Lewkowycz et al., 2022](#)). Attention-based models tend to excel best in tasks involving language and when trained on massively large corpora; their ability to extrapolate in manipulation tasks ([Zhou et al., 2022](#)) and in more modest data regimes remains an area of ongoing research. Moreover, while recent research has aimed to study their (in-distribution) generalization properties ([Edelman et al., 2022](#)), their capacity for “reasoning” more generally remains mysterious ([Zhang et al., 2022b](#)).

Lastly, a line of research has studied various bilinear models (e.g. [Hong et al. \(2021\)](#); [Shah et al. \(2020\)](#)), motivated by the extension of literature on matrix factorization discussed in our abridged related work. However, these methods require additional fine tuning when applied to novel goals and do not display zero shot extrapolation. Moreover, [Shah et al. \(2020\)](#) requires a discretization of the state-action space in order to formulate the matrix completion problem, whereas bilinear transduction shows how the results hold in continuous bilinear form, without the need for discretization. Along similar lines, the Deep-Sets architecture of [Zaheer et al. \(2017\)](#) aims for combinatorial extrapolation by embedding tokens of interest in a latent vector space on which addition operators can be defined. [Zhou et al. \(2022\)](#) compares the Deep Sets and Transformers approaches as policy architectures in reinforcement learning, finding that neither architecture uniformly outperforms the other.

Distributional Robustness. One popular approach to out-of-distribution learning has been distributional robustness, which seeks predictors that perform on a family of “nearby” shifted distributions

Sinha et al. (2018); Rahimian & Mehrotra (2019). These approaches are well suited to OOD settings where the test distributions have the same support, but have differing (but boundedly different) probability densities.

Meta-Learning and Multi-Task Learning: Meta-learning and multi-task learning methods aim to learn policies/predictors that can either zero-shot or very quickly generalize to new problems Thrun & Pratt (1998); Finn et al. (2017a); Santoro et al. (2016); Vinyals et al. (2016); Kalashnikov et al. (2021); Crawshaw (2020); Rosenbaum et al. (2018); Caruana (1997). These methods however usually make strict distributional assumptions about the training and test distribution of tasks being the same. So while the particular data distribution may be different, the meta-level assumption is still an in-distribution one Fallah et al. (2020). Yin et al. (2019) shows that under distribution shift, meta-learning methods can fail dramatically, and this problem will be exacerbated when supports shift as well. Multi-task learning methods often make the assumption that training on some set of training tasks can provide good pre-training before finetuning on new tasks Julian et al. (2020); Su et al. (2022); Meftah et al. (2020). However, it is poorly understood how these tasks actually relate to each other and the study has been largely empirical.

Generalization in Reinforcement Learning and Imitation Learning: Reinforcement learning problems are unique in that they are sequential and interactive. Distribution shift may not just be across different MDPs but within the states of a single MDP as well Filos et al. (2020); Kirk et al. (2021); Ghosh et al. (2021). While we do not explicitly focus on the reinforcement learning setting, our results are largely in the imitation learning setting. Moreover, we are not looking at the finetuning setting but rather evaluating zero-shot performance on out of support goals, dynamics and arbitrary contexts. While certain works Simeonov et al. (2022); Wang et al. (2022) do explore ideas of equivariance and invariance in RL and imitation learning, they are largely restricted to the either the $SE(3)$ or $SO(2)$ groups. Moreover, as we discussed in our experiments, bilinear transduction is able to capture more complex notions than equivariance.

The empirical evaluations we perform are in a large part on domains in goal conditioned reinforcement learning. Most goal conditioned RL works do not consider extrapolation beyond the range of training goals, but rather standard statistical generalization amongst training goals Kaelbling (1993); Andrychowicz et al. (2017); Nair et al. (2018). While certain architectures Hong et al. (2021) do seem to show some level of extrapolation, they still require significant finetuning. We show that with low-rank assumptions under reparameterization, zero-shot extrapolation is possible.

Our work can also loosely be connected to analogy making techniques in machine learning Reed et al. (2015); Ichien et al. (2021); Hansen-Estruch et al. (2022); Mitchell (2021). We formalize this notion in the context of out-of-support extrapolation and show that transduction allows for analogy making under low-rank structure assumptions.

B THEORETICAL RESULTS

This appendix is organized into four parts.

- Appendix B.1 introduces the bounded density condition, and reproduces a folklore guarantee for extrapolation when one distribution has bounded density with respect to the other.
- Appendix B.2 provides a basic guarantee for extrapolation in the context of missing-not-at-random matrix completion depicted in Fig 9, based on results in Shah et al. (2020).
- Appendix B.3 formalizes a notion of bounded combinatorial density ratio (Definition B.2), in terms of which we can establish out-of-combination extrapolation guarantees by leveraging Appendix B.2.
- Finally, Appendix B.4 applies the results of the previous section, stating and proving formal guarantees for our proposed transductive predictors.

B.1 GENERALIZATION UNDER BOUNDED DENSITY RATIO

The following gives a robust, quantitative notion of when one distribution is in the support of another. For generality, we state this condition in terms of general positive measures μ_1, μ_2 , which need not be normalized and sum to one.

Definition B.1 (κ -bounded density ratio). Let μ_1, μ_2 be two measures over a space Ω . We say μ_1 has κ -bounded density with respect to μ_2 , which we denote $\mu_1 \ll_{\kappa} \mu_2$, if for all measurable event³ $A \subset \Omega$, $\mu_1[A] \leq \kappa \mu_2[A]$.

Stating Definition B.1 for general probability affords us the flexibility to write example, $\mathbb{P}_1 \ll_{\kappa} \mathbb{P}_2 + \mathbb{P}_3$, as $\mathbb{P}_2 + \mathbb{P}_3$ is a nonnegative measure with total mass $1 + 1 = 2$.

Remark B.1 (Connection to concentrability). The parameter κ is known in the off-policy reinforcement learning literature as the *concentrability* coefficient (see, e.g. Munos & Szepesvári (2008)), and appears in controlling the performance of a rollout policy π_1 trained on a behavior policy π_2 by asserting that $\mathbb{P}_{\pi_1} \ll_{\kappa} \mathbb{P}_{\pi_2}$, where \mathbb{P}_{π_i} is, for example, a state-action visitation probability under π_i .

Remark B.2 (Density nomenclature). The nomenclature “density” refers to the alternative definition in terms of Radon-Nikodym derivatives. To avoid technicalities, this can be best seen when μ_1 and μ_2 are continuous densities over $\Omega = \mathbb{R}^d$ with densities $p_1(\cdot)$ and $p_2(\cdot)$; e.g. for $i \in \{1, 2\}$, $\mu_i[A] = \int_{x \in A} p_i(x) dx$. Then $\mu_1 \ll_{\kappa} \mu_2$ is equivalent to

$$\sup_x \frac{p_1(x)}{p_2(x)} \leq \kappa.$$

The following lemma motivates the use of Definition B.1. Its proof is standard but included for completeness.

Lemma B.1. Let μ_1, μ_2 be measures on the same measurable space Ω , and suppose that $\mu_2 \ll_{\kappa} \mu_1$. Then, for any nonnegative function ϕ , $\mu_2[\phi] \leq \kappa \mu_1[\phi]$.⁴ In particular, if $\mathcal{D}_{\text{test}} \ll_{\kappa} \mathcal{D}_{\text{train}}$, then as long as our loss function $\ell(\cdot, \cdot)$ is nonnegative,

$$\mathcal{R}(h_{\theta}; \mathcal{D}_{\text{test}}) \leq \kappa \mathcal{R}(h_{\theta}; \mathcal{D}_{\text{train}}).$$

Thus, up to a κ -factor, $\mathcal{R}(h_{\theta}; \mathcal{D}_{\text{test}})$ inherits any in-distribution generalization guarantees for $\mathcal{R}(h_{\theta}; \mathcal{D}_{\text{train}})$.

Proof. As in standard measure theory (c.f. Çinlar (2011, Chapter 1)), we can approximate any $\phi \geq 0$ by a sequence of simple functions $\phi_n \uparrow \phi$, where $\phi_n(\omega) = \sum_{i=1}^{k_n} c_{n,i} \mathbb{I}\{\omega \in A_{n,i}\}$, with $A_{n,i} \subset \Omega$ and $c_{n,i} \geq 0$. For each ϕ_n , we have

$$\mu_2[\phi_n] = \sum_{i=1}^{k_n} c_{n,i} \mu_2[A_{n,i}] \leq \kappa \sum_{i=1}^{k_n} c_{n,i} \mu_1[A_{n,i}] = \mu_1[\phi_n].$$

The result now follows from the monotone convergence theorem. To derive the special case for $\mathcal{D}_{\text{test}}$ and $\mathcal{D}_{\text{train}}$, apply the general result with nonnegative function $\phi(x) = \mathbb{E}_{y \sim h_{\theta}(x)} \ell(y, h_{\star}(x))$ (recall $\ell(\cdot, \cdot) \geq 0$ by assumption), $\mu_1 = \mathcal{D}_{\text{train}}$ and $\mu_2 = \mathcal{D}_{\text{test}}$. \square

³For simplicity, we omit concrete discussion of measurability concerns throughout.

⁴Here, $\mu[\phi] := \int \phi(\omega) d\mu(\omega)$ denotes the integration with respect to μ .

B.2 EXTRAPOLATION FOR MATRIX COMPLETION

In what follows, we derive a simple extrapolation guarantee for matrix completion. The following is in the spirit of the Nyström column approximation (see e.g. [Gittens & Mahoney \(2013\)](#)), and our proof follows the analysis due to [Shah et al. \(2020\)](#). Throughout, consider

$$\hat{\mathbf{M}} = \begin{bmatrix} \hat{\mathbf{M}}_{11} & \hat{\mathbf{M}}_{12} \\ \hat{\mathbf{M}}_{21} & \hat{\mathbf{M}}_{22} \end{bmatrix}, \quad \mathbf{M}^* = \begin{bmatrix} \mathbf{M}_{11}^* & \mathbf{M}_{12}^* \\ \mathbf{M}_{21}^* & \mathbf{M}_{22}^* \end{bmatrix},$$

where we decompose $\hat{\mathbf{M}}, \mathbf{M}^*$ into blocks $(i, j) \in \{1, 2\}^2$ for dimension $n_i \times m_j$.

Lemma B.2. *Suppose that $\hat{\mathbf{M}}$ is rank at most p , \mathbf{M}^* is rank p , and*

$$\forall (i, j) \neq (2, 2), \quad \|\hat{\mathbf{M}}_{i,j} - \mathbf{M}_{i,j}^*\|_{\text{F}} \leq \epsilon, \quad \text{and} \quad \|\mathbf{M}_{i,j}^*\|_{\text{F}} \leq M,$$

where $\epsilon \leq \sigma_p(\mathbf{M}_{11}^*)/2$. Then,

$$\|\hat{\mathbf{M}}_{22} - \mathbf{M}_{22}^*\|_{\text{F}} \leq 8\epsilon \frac{M^2}{\sigma_p(\mathbf{M}_{11}^*)^2}.$$

Proof. The proof mirrors that of [Shah et al. \(2020, Proposition 13\)](#). We shall show below that $\hat{\mathbf{M}}$ is of rank exactly p . Hence, [Shah et al. \(2020, Lemma 12\)](#) gives the following exact expression for the bottom-right blocks,

$$\hat{\mathbf{M}}_{22} = \hat{\mathbf{M}}_{21} \hat{\mathbf{M}}_{11}^\dagger \hat{\mathbf{M}}_{12}, \quad \mathbf{M}_{22}^* = \mathbf{M}_{21}^* (\mathbf{M}_{11}^*)^\dagger \mathbf{M}_{12}^*,$$

where above $(\cdot)^\dagger$ denotes the Moore-Penrose pseudoinverse. Since $\|\hat{\mathbf{M}}_{11} - \mathbf{M}_{11}^*\|_{\text{op}} \leq \|\hat{\mathbf{M}}_{11} - \mathbf{M}_{11}^*\|_{\text{F}} \leq \epsilon \leq \sigma_p(\mathbf{M}_{11}^*)/2$, Weyl's inequality implies that $\hat{\mathbf{M}}_{11}$ is rank p (as promised), and $\|\hat{\mathbf{M}}_{11}^\dagger\|_{\text{op}} \leq 2\sigma_p(\mathbf{M}_{11}^*)^{-1}$. Similarly, as $\|\hat{\mathbf{M}}_{12} - \mathbf{M}_{12}^*\|_{\text{op}} \leq \sigma_p(\mathbf{M}_{11}^*)/2 \leq M/2$, so $\|\hat{\mathbf{M}}_{12}\|_{\text{op}} \leq \frac{3}{2}M$. Thus,

$$\begin{aligned} \|\hat{\mathbf{M}}_{22} - \mathbf{M}_{22}^*\|_{\text{F}} &\leq \|\hat{\mathbf{M}}_{21} - \mathbf{M}_{21}^*\|_{\text{F}} \|\hat{\mathbf{M}}_{11}^\dagger\|_{\text{op}} \|\hat{\mathbf{M}}_{12}\|_{\text{op}} + \|\mathbf{M}_{21}^*\|_{\text{op}} \|\hat{\mathbf{M}}_{11}^\dagger\|_{\text{op}} \|\mathbf{M}_{12}^* - \hat{\mathbf{M}}_{12}\|_{\text{F}} \\ &\quad + \|\mathbf{M}_{12}^*\|_{\text{op}} \|\mathbf{M}_{21}^*\|_{\text{op}} \|\hat{\mathbf{M}}_{11}^\dagger - (\mathbf{M}_{11}^*)^\dagger\|_{\text{F}} \\ &\leq \frac{5\epsilon M}{2\sigma_p(\mathbf{M}^*)} + M^2 \|\hat{\mathbf{M}}_{11}^\dagger - (\mathbf{M}_{11}^*)^\dagger\|_{\text{F}}. \end{aligned}$$

Next, using a perturbation bound on the pseudoinverse⁵ due to [Meng & Zheng \(2010, Theorem 2.1\)](#),

$$\begin{aligned} \|\hat{\mathbf{M}}_{11}^\dagger - (\mathbf{M}_{11}^*)^\dagger\|_{\text{F}} &\leq \|\hat{\mathbf{M}}_{11} - \mathbf{M}_{11}^*\|_{\text{F}} \max\{\|\hat{\mathbf{M}}_{11}^\dagger\|_{\text{op}}^2, \|(\mathbf{M}_{11}^*)^\dagger\|_{\text{op}}^2\} \\ &\leq \epsilon \cdot 4\sigma_p(\mathbf{M}_{11}^*)^{-2}. \end{aligned}$$

Therefore, we conclude

$$\|\hat{\mathbf{M}}_{22} - \mathbf{M}_{22}^*\|_{\text{F}} \leq \frac{5\epsilon M}{2\sigma_p(\mathbf{M}^*)} + \epsilon \frac{4M^2}{\sigma_p(\mathbf{M}_{11}^*)^2} \leq 8\epsilon \frac{M^2}{\sigma_p(\mathbf{M}_{11}^*)^2}. \quad \square$$

B.3 GENERAL ANALYSIS FOR COMBINATORIAL EXTRAPOLATION

We now provide our general analysis for combinatorial extrapolation. To avoid excessive subscripts, we write $\mathcal{X} = \mathcal{W} \times \mathcal{V}$ rather than $\mathcal{X} = \mathcal{X}_1 \times \mathcal{X}_2$ as in the main body. We consider extrapolation under the following definition of combinatorial support.

Definition B.2 (Bounded combinatorial density ratio). Let $\mathcal{D}, \mathcal{D}'$ be two distributions over a product space $\mathcal{W} \times \mathcal{V}$. We say \mathcal{D}' has κ -bounded combinatorial density ratio with respect to \mathcal{D} , written as $\mathcal{D}' \ll_{\kappa, \text{comb}} \mathcal{D}$, if there exist distributions $\mathcal{D}_{\mathcal{W},i}$ and $\mathcal{D}_{\mathcal{V},j}$, $i, j \in \{1, 2\}$, over \mathcal{W} and \mathcal{V} , respectively, such that $\mathcal{D}_{i \otimes j} := \mathcal{D}_{\mathcal{W},i} \otimes \mathcal{D}_{\mathcal{V},j}$ satisfy

$$\sum_{(i,j) \neq (2,2)} \mathcal{D}_{i \otimes j} \ll_{\kappa} \mathcal{D}, \quad \text{and} \quad \mathcal{D}' \ll_{\kappa} \sum_{i,j=1,2} \mathcal{D}_{i \otimes j}.$$

⁵Unlike [Shah et al. \(2020\)](#), we are interested in the Frobenius norm error, so we elect for the slightly sharper bound of [Meng & Zheng \(2010\)](#) above than the classical operator norm bound of [Stewart \(1977\)](#).

For simplicity, we consider scalar predictors, as the general result for vector valued estimators can be obtained by stacking the components. Specifically, we consider a ground-truth predictor h_* and estimator \hat{h} of the form

$$h_* = \langle f_*, g_* \rangle, \quad \hat{h} = \langle \hat{f}, \hat{g} \rangle, \quad f_*, \hat{f} : \mathcal{W} \rightarrow \mathbb{R}^p, \quad g_*, \hat{g} : \mathcal{V} \rightarrow \mathbb{R}^p. \quad (\text{B.1})$$

Lastly, we choose the (scalar) square-loss, yielding the following risk

$$\mathcal{R}(\hat{h}; \mathcal{D}) := \mathbb{E}_{(w,v) \sim \mathcal{D}} [(h_*(w, v) - \hat{h}(w, v))^2].$$

Throughout, we assume that all expectations that arise are finite. Our main guarantee is as follows.

Theorem 2. *Let $\mathcal{D}, \mathcal{D}'$ be two distributions on $\mathcal{W} \times \mathcal{V}$ satisfying $\mathcal{D}' \ll_{\kappa, \text{comb}} \mathcal{D}$, with corresponding factor distributions $\mathcal{D}_{\mathcal{W}, i}$ and $\mathcal{D}_{\mathcal{V}, j}$, $i, j \in \{1, 2\}$. Define the effective singular value*

$$\sigma_*^2 := \sigma_p(\mathbb{E}_{\mathcal{D}_{\mathcal{W}, 1}} [f_*(w) f_*(w)^\top]) \sigma_p(\mathbb{E}_{\mathcal{D}_{\mathcal{V}, 1}} [g_*(v) g_*(v)^\top]), \quad (\text{B.2})$$

and suppose that $\max_{1 \leq i, j \leq 2} \mathbb{E}_{\mathcal{D}_{i \otimes j}} |h_*(w, v)|^2 \leq M_*^2$. $\mathcal{R}(\hat{h}; \mathcal{D}) \leq \frac{\sigma_*^2}{4\kappa}$,

$$\mathcal{R}(\hat{h}; \mathcal{D}') \leq \mathcal{R}(\hat{h}; \mathcal{D}) \cdot \kappa^2 \left(1 + 64 \frac{M_*^4}{\sigma_*^4} \right) = \mathcal{R}(\hat{h}; \mathcal{D}) \cdot \text{poly} \left(\kappa, \frac{M_*}{\sigma_*} \right).$$

B.3.1 PROOF OF THEOREM 2

First, let us assume the following two conditions hold; we shall derive these conditions from the conditions of [Theorem 2](#) at the end of the proof:⁶

$$\forall (i, j) \neq (2, 2), \quad \mathcal{R}(\hat{h}; \mathcal{D}_{i \otimes j}) \leq \epsilon^2, \quad \mathbb{E}_{\mathcal{D}_{i \otimes j}} [h_*(w, v)^2] \leq M_*^2, \quad \epsilon < \sigma_*/2. \quad (\text{B.3})$$

Our strategy is first to prove a version of [Theorem 2](#) for when \mathcal{W} and \mathcal{V} have finite cardinality by reduction to the analysis of matrix completion in [Lemma B.2](#), and then extend to arbitrary domains via a limiting argument.

Lemma B.3. *Suppose that [Eq. \(B.3\)](#) hold, and in addition, that \mathcal{W} and \mathcal{V} have finite cardinality. Then,*

$$\mathcal{R}(\hat{h}; \mathcal{D}_{2 \otimes 2}) = \|\hat{\mathbf{M}}_{22} - \mathbf{M}_{22}^*\|_{\text{F}}^2 \leq 64\epsilon^2 \frac{M_*^4}{\sigma_*^4}.$$

Proof [Lemma B.3](#). By adding additional null elements to either \mathcal{W} or \mathcal{V} , we may assume without loss of generality that $|\mathcal{W}| = |\mathcal{V}| = d$, and enumerate their elements $\{w_1, \dots, w_d\}$ and $\{v_1, \dots, v_d\}$. Let $\mathbf{p}_{i,a} = \Pr_{w \sim \mathcal{D}_{\mathcal{W}, i}} [w = w_a]$ and $\mathbf{q}_{j,b} = \Pr_{v \sim \mathcal{D}_{\mathcal{V}, j}} [v = v_b]$. Consider matrices $\hat{\mathbf{M}}, \mathbf{M}^* \in \mathbb{R}^{2d \times 2d}$, with $d \times d$ blocks

$$(\hat{\mathbf{M}}_{ij})_{ab} = \sqrt{\mathbf{p}_{i,a} \mathbf{q}_{j,b}} \cdot \hat{h}(w_a, v_b), \quad (\mathbf{M}_{ij}^*)_{ab} = \sqrt{\mathbf{p}_{i,a} \mathbf{q}_{j,b}} \cdot h_*(w_a, v_b).$$

We then verify that

$$\begin{aligned} \|\hat{\mathbf{M}}_{ij} - \mathbf{M}_{ij}^*\|_{\text{F}}^2 &= \sum_{a,b=1}^d \mathbf{p}_{i,a} \mathbf{q}_{j,b} (\hat{h}(w_a, v_b) - h_*(w_a, v_b))^2 \\ &= \mathbb{E}_{\mathcal{D}_{i \otimes j}} [(\hat{h}(w, v) - h_*(w, v))^2] = \mathcal{R}(\hat{h}; \mathcal{D}_{i \otimes j}), \end{aligned} \quad (\text{B.4})$$

and thus $\|\hat{\mathbf{M}}_{ij} - \mathbf{M}_{ij}^*\|_{\text{F}}^2 \leq \epsilon^2$ for $(i, j) \neq (2, 2)$. Furthermore, define the matrices $\hat{\mathbf{A}}_i, \hat{\mathbf{B}}_j$ via

$$(\hat{\mathbf{A}}_i)_a := \sqrt{\mathbf{p}_{i,a}} \hat{f}(w_a)^\top, \quad (\hat{\mathbf{B}}_j)_b := \sqrt{\mathbf{q}_{j,b}} \hat{g}(v_b)^\top,$$

and define $\mathbf{A}_i^*, \mathbf{B}_j^*$ similarly. Then,

$$\hat{\mathbf{M}} = \begin{bmatrix} \hat{\mathbf{A}}_1 \\ \hat{\mathbf{A}}_2 \end{bmatrix} \begin{bmatrix} \hat{\mathbf{B}}_1 \\ \hat{\mathbf{B}}_2 \end{bmatrix}^\top, \quad \mathbf{M}^* = \begin{bmatrix} \mathbf{A}_1^* \\ \mathbf{A}_2^* \end{bmatrix} \begin{bmatrix} \mathbf{B}_1^* \\ \mathbf{B}_2^* \end{bmatrix}^\top,$$

⁶Notice that here we take M_*^2 as an upper bound of $\mathbb{E}_{\mathcal{D}_{i \otimes j}} [h_*(w, v)^2]$, rather than a pointwise upper bound in [Theorem 2](#). This is for convenience in a limiting argument below.

showing that $\text{rank}(\hat{\mathbf{M}}_1), \text{rank}(\hat{\mathbf{M}}_2) \leq p$. Finally, by Eq. (B.2),

$$\begin{aligned} \sigma_p(\mathbf{M}_{11}^*)^2 &= \sigma_p(\mathbf{A}_1^*(\mathbf{B}_1^*)^\top)^2 \geq \sigma_p^2(\mathbf{A}_1^*)\sigma_p^2(\mathbf{B}_1^*) \\ &= \sigma_p\left(\left(\mathbf{A}_1^*\right)^\top \mathbf{A}_1^*\right) \sigma_p\left(\left(\mathbf{B}_1^*\right)^\top \mathbf{B}_1^*\right) \\ &= \sigma_p\left(\sum_{a=1}^d \mathfrak{p}_{1,a} \hat{f}(w_a) \hat{f}(w_a)^\top\right) \sigma_p\left(\sum_{b=1}^d \mathfrak{q}_{1,b} \hat{g}(v_b) \hat{g}(v_b)^\top\right) \\ &= \sigma_p(\mathbb{E}_{\mathcal{D}_{\mathcal{W},1}}[\hat{f}(w) \hat{f}(w)^\top]) \sigma_p(\mathbb{E}_{\mathcal{D}_{\mathcal{V},1}}[\hat{g}(v) \hat{g}(v)^\top]) = \sigma_\star^2. \end{aligned}$$

Lastly, we have

$$\|\mathbf{M}_{i,j}^*\|_{\mathbb{F}}^2 = \sum_{a,b} \mathfrak{p}_{i,a} \mathfrak{q}_{j,b} h_\star(w_a, v_b)^2 = \mathbb{E}_{\mathcal{D}_{i \otimes j}} h_\star(w, v)^2 \leq M_\star^2.$$

Thus, Eq. (B.4) and Lemma B.2 imply that

$$\mathcal{R}(\hat{h}; \mathcal{D}_{2 \otimes 2}) = \|\hat{\mathbf{M}}_{22} - \mathbf{M}_{22}^*\|_{\mathbb{F}}^2 \leq 64\epsilon^2 \frac{M_\star^4}{\sigma_\star^4}.$$

□

Lemma B.4. Suppose that Eq. (B.3) hold, but unlike Lemma B.4, \mathcal{W} and \mathcal{V} need not be finite spaces. Then, still, it holds that

$$\mathcal{R}(\hat{h}; \mathcal{D}_{2 \otimes 2}) = \|\hat{\mathbf{M}}_{22} - \mathbf{M}_{22}^*\|_{\mathbb{F}}^2 \leq 64\epsilon^2 \frac{M_\star^4}{\sigma_\star^4}.$$

Proof of Lemma B.4. For $n \in \mathbb{N}$, define $h_{\star,n} = \langle f_{\star,n}, g_{\star,n} \rangle$ and $\hat{h}_n = \langle \hat{f}_n, \hat{g}_n \rangle$, where $f_{\star,n}, \hat{f}_n, \hat{g}_n, g_{\star,n}$ are simple functions (i.e. finite range, see the proof of Lemma B.1) converging to $f_\star, \hat{f}, g_\star, \hat{g}$. Define

$$\begin{aligned} \sigma_{\star,n}^2 &= \sigma_p(\mathbb{E}_{\mathcal{D}_{\mathcal{W},1}}[f_{\star,n}(w) f_{\star,n}(w)^\top]) \sigma_p(\mathbb{E}_{\mathcal{D}_{\mathcal{V},1}}[g_{\star,n}(v) g_{\star,n}(v)^\top]), \\ M_{\star,n}^2 &= \max_{i,j \neq (2,2)} \mathbb{E}_{\mathcal{D}_{i \otimes j}} [h_{\star,n}(w, v)^2] \\ \epsilon_n^2 &= \max_{i,j \neq (2,2)} \mathcal{R}(\hat{h}_n; \mathcal{D}_{i \otimes j}). \end{aligned}$$

By the dominated convergence theorem⁷,

$$\liminf_{n \geq 1} \sigma_{\star,n}^2 \geq \sigma_\star^2, \quad \limsup_{n \geq 1} M_{\star,n}^2 \leq M_\star^2, \quad \limsup_{n \geq 1} \epsilon_n^2 \leq \epsilon^2.$$

In particular, as $\epsilon^2 \leq \sigma^2/4$, then applying Lemma B.3 for n sufficiently large,

$$\mathcal{R}(\hat{h}_n; \mathcal{D}_{2 \otimes 2}) \leq 64 \frac{M_{\star,n}^4}{\sigma_{\star,n}^4} \epsilon_n^2.$$

Indeed, for any fixed n , all of $\hat{f}_n, \hat{g}_n, f_{\star,n}, g_{\star,n}$ are simple functions, so we can partition \mathcal{W} and \mathcal{V} into sets on which these embeddings are constant, and thus treat \mathcal{W} and \mathcal{V} as finite domains; this enables the application of Lemma B.3 applies. Finally, using the dominated coverage theorem one last time,

$$\mathcal{R}(\hat{h}; \mathcal{D}_{2 \otimes 2}) = \lim_{n \rightarrow \infty} \mathcal{R}(\hat{h}_n; \mathcal{D}_{2 \otimes 2}) \leq \limsup_{n \geq 1} 64 \frac{M_{\star,n}^4}{\sigma_{\star,n}^4} \epsilon_n^2 \leq 64 \frac{M_\star^4}{\sigma_\star^4} \epsilon^2.$$

□

We can now conclude the proof of our proposition.

⁷Via standard arguments, one can construct the limiting embeddings $f_\star, \hat{f}_n, \hat{g}_n, g_\star$ in such a way that their norms are dominated by integrable functions.

Proof of Theorem 2. As $\mathcal{D}' \ll_{\kappa} \sum_{i,j} \mathcal{D}_{i \otimes j}$ and $\sum_{(i,j) \neq (2,2)} \mathcal{D}_{i \otimes j} \ll_{\kappa} \mathcal{D}$, Lemma B.1 and additivity of the integral implies

$$\begin{aligned} \mathcal{R}(\hat{h}; \mathcal{D}') &\leq \kappa \mathcal{R}(\hat{h}; \mathcal{D}_{2 \otimes 2}) + \kappa \sum_{(i,j) \neq (2,2)} \mathcal{R}(\hat{h}; \mathcal{D}_{i \otimes j}) \\ &\leq \kappa \mathcal{R}(\hat{h}; \mathcal{D}_{2 \otimes 2}) + \kappa^2 \mathcal{R}(\hat{h}; \mathcal{D}). \end{aligned} \quad (\text{B.5})$$

Moreover, setting $\epsilon^2 := \kappa \mathcal{R}(\hat{h}; \mathcal{D})$, we have

$$\max_{(i,j) \neq (2,2)} \mathcal{R}(\hat{h}; \mathcal{D}_{i \otimes j}) \leq \sum_{(i,j) \neq (2,2)} \mathcal{R}(\hat{h}; \mathcal{D}_{i \otimes j}) \leq \kappa \mathcal{R}(\hat{h}; \mathcal{D}) := \epsilon^2.$$

Thus, for $\mathcal{R}(\hat{h}; \mathcal{D}) < \frac{\sigma_{\star}^2}{4\kappa}$, Eq. (B.3) holds and thus Lemma B.4 entails

$$\mathcal{R}(\hat{h}; \mathcal{D}_{2 \otimes 2}) \leq 64\epsilon^2 \frac{M_{\star}^4}{\sigma_{\star}^4} = 64\kappa \mathcal{R}(\hat{h}; \mathcal{D}) \frac{M_{\star}^4}{\sigma_{\star}^4}.$$

Thus, combining with Eq. (B.5),

$$\mathcal{R}(\hat{h}; \mathcal{D}') \leq \kappa^2 \mathcal{R}(\hat{h}; \mathcal{D}) \cdot \left(1 + 64 \frac{M_{\star}^4}{\sigma_{\star}^4}\right),$$

completing the proof. \square

B.4 EXTRAPOLATION FOR TRANSDUCTION

Leveraging Theorem 2, this section proves Theorem 1, thereby providing a formal theoretical justification for predictors of the form Eq. (3.5).

Proof of Theorem 1. We argue by reducing to Theorem 2. The reparameterization of the stochastic predictor h_{θ} in Eq. (3.1), followed by Assumption 3.2 allows us to write

$$\begin{aligned} \mathcal{R}(h_{\theta}; \mathcal{D}_{\text{train}}) &= \mathbb{E}_{x \sim \mathcal{D}_{\text{train}}} \mathbb{E}_{y \sim h_{\theta}(x)} \ell(y, h_{\star}(x)) \\ &= \mathbb{E}_{x \sim \mathcal{D}_{\text{train}}} \mathbb{E}_{x' \sim \mathcal{D}_{\text{trns}}(x)} \ell(\bar{h}_{\theta}(x - x', x'), h_{\star}(x)) \\ &= \mathbb{E}_{x \sim \mathcal{D}_{\text{train}}} \mathbb{E}_{x' \sim \mathcal{D}_{\text{trns}}(x)} \ell(\bar{h}_{\theta}(x - x', x'), \bar{h}_{\star}(x - x', x')). \end{aligned}$$

In the above display, the joint distribution of $(x - x', x')$ is precisely given by $\bar{\mathcal{D}}_{\text{train}}$ (see Eq. (3.2)). Hence,

$$\mathcal{R}(h_{\theta}; \mathcal{D}_{\text{train}}) = \mathbb{E}_{\bar{\mathcal{D}}_{\text{train}}} \ell(\bar{h}_{\theta}(\Delta x, x'), \bar{h}_{\star}(\Delta x, x')).$$

Further, as $\ell(y, y') = \|y - y'\|^2$ is the square loss and decomposes across coordinates,

$$\mathcal{R}(h_{\theta}; \mathcal{D}_{\text{train}}) = \sum_{k=1}^K \mathbb{E}_{\bar{\mathcal{D}}_{\text{train}}} (\bar{h}_{\theta,k}(\Delta x, x') - \bar{h}_{\star,k}(\Delta x, x'))^2. \quad (\text{B.6})$$

By the same token,

$$\mathcal{R}(h_{\theta}; \mathcal{D}_{\text{test}}) = \sum_{k=1}^K \mathbb{E}_{\bar{\mathcal{D}}_{\text{test}}} (\bar{h}_{\theta,k}(\Delta x, x') - \bar{h}_{\star,k}(\Delta x, x'))^2.$$

To conclude the proof, we remain to show that for all $k \in [K]$, we have

$$\begin{aligned} \mathbb{E}_{\bar{\mathcal{D}}_{\text{test}}} (\bar{h}_{\theta,k}(\Delta x, x') - \bar{h}_{\star,k}(\Delta x, x'))^2 &\leq C_{\text{prob}} \cdot \mathbb{E}_{\bar{\mathcal{D}}_{\text{train}}} (\bar{h}_{\theta,k}(\Delta x, x') - \bar{h}_{\star,k}(\Delta x, x'))^2, \\ \text{where } C_{\text{prob}} &= \kappa^2 \left(1 + 64 \frac{M_{\star}^4}{\sigma_{\star}^4}\right). \end{aligned} \quad (\text{B.7})$$

Indeed, for each $k \in [K]$, we have

$$\mathbb{E}_{\bar{\mathcal{D}}_{\text{train}}} (\bar{h}_{\theta,k}(\Delta x, x') - \bar{h}_{\star,k}(\Delta x, x'))^2 \stackrel{(\text{Eq. (B.6)})}{\leq} \mathcal{R}(h_{\theta}; \mathcal{D}_{\text{train}}) \stackrel{(\text{by assumption})}{\leq} \frac{\sigma^2}{4\kappa}.$$

Hence Eq. (B.7) holds by invoking Theorem 2 with the correspondences $\mathcal{W} \leftarrow \Delta \mathcal{X}$, $\mathcal{V} \leftarrow \mathcal{X}$, $\sigma_{\star} \leftarrow \sigma$, $M_{\star} \leftarrow M$ and $\kappa \leftarrow \kappa$. This concludes the proof. \square

B.4.1 FURTHER REMARKS ON THEOREM 1

- The singular value condition, $\min\{\sigma_p(\mathbb{E}_{\mathcal{D}_{\Delta x,1}}[f_{*,k}f_{*,k}^\top]), \sigma_p(\mathbb{E}_{\mathcal{D}_{x,1}}[g_{*,k}g_{*,k}^\top])\} \geq \sigma^2 > 0$, mirrors the non-degeneracy conditions given in the past work in matrix completion (c.f. (Shah et al., 2020)).
- The support condition $\sup_{\Delta x, x'} |\bar{h}_{*,k}(\Delta x, x')| \leq M$ is a mild boundedness condition, which (in light of Theorem 2) can be weakened further to

$$\max_{1 \leq i, j \leq 2} \mathbb{E}_{\mathcal{D}_{i \otimes j}} [\bar{h}_{*,k}(\Delta x, x')^2] \leq M^2,$$

where $\mathcal{D}_{i \otimes j}$ are the constituent distributions witnessing $\bar{\mathcal{D}}_{\text{test}} \ll_{\kappa, \text{comb}} \bar{\mathcal{D}}_{\text{train}}$.

- The final condition, $\mathcal{R}(h_\theta; \mathcal{D}_{\text{train}}) \leq \frac{\sigma^2}{4\kappa}$, is mostly for convenience. Indeed, as $M \geq \sigma$ and $\kappa \geq 1$, then as soon as $\mathcal{R}(h_\theta; \mathcal{D}_{\text{train}}) > \frac{\sigma^2}{4\kappa}$, our upper-bound on $\mathcal{R}(h_\theta; \mathcal{D}_{\text{test}})$ is no better than

$$\kappa M^2 \cdot \frac{64}{4} \cdot \frac{M^2}{\sigma^2} \geq 6M^2,$$

which is essentially vacuous. Indeed, if we also inflate M and stipulate that $\sup_{\Delta x, x'} |\bar{h}_\theta(\Delta x, x')| \leq \sqrt{6}M$, we can remove the condition $\mathcal{R}(h_\theta; \mathcal{D}_{\text{train}}) \leq \frac{\sigma^2}{4\kappa}$ altogether.

B.5 CONNECTIONS TO MATRIX COMPLETION




Figure 9: Illustration of bilinear representations for OOC learning, and connection to matrix completion. **(a)** An example of low-rank matrix completion, where both M and M_{11} have rank- p . Blue: support where entries can be accessed, green: entries are missing. **(b)** An example that low-rank structure facilitates certain forms of OOC, i.e. for each $k \in [K]$, the predictor can be represented by bilinear embeddings as $\bar{h}_{\theta,k}(\Delta x, x') = \langle f_{\theta,k}(\Delta x), g_{\theta,k}(x') \rangle$.

Building upon Definition B.1, Assumption 3.1 introduces a notion of bounded density ratio between $\bar{\mathcal{D}}_{\text{train}}$ and $\bar{\mathcal{D}}_{\text{test}}$ in the OOC setting. Take the discrete case of matrix completion as an example, as illustrated in Fig 9, the training distribution of $(\Delta x, x')$ covers the support of the (1, 1), (1, 2), (2, 1) blocks of the matrix, while the test distribution of $(\Delta x, x')$ might be covered by any product of the marginals of the 2×2 blocks. With this connection in mind, it is possible to establish the OOC guarantees on $\bar{\mathcal{D}}_{\text{test}}$ as in matrix completion tasks, if the bilinear embedding admits some low-rank structure. In other words, samples from the top-left and off-diagonal blocks uniquely determine the bottom-right block.

This is a standard setting studied extensively in the matrix completion literature, see (Shah et al., 2020; Agarwal et al., 2021b; Athey et al., 2021) for example. Our results in Theorem 1 can be viewed as a generalization of these results to the continuous space embeddings. Finally, note that Assumptions 3.2 and 3.3 are also adopted correspondingly from the matrix completion literature, where Assumption 3.2 concerns the ground-truth bilinear structure of the embeddings, which mirrors the *low-rank* structure of the underlying matrix in the matrix completion case; Assumption 3.3 corresponds to the non-degeneracy assumption on the the top-left block distribution, as referred to as “anchor” columns/rows in the past work when dealing with matrices (Shah et al., 2020).

C ADDITIONAL RESULTS

C.1 LEARNING WHICH POINTS TO TRANSDUCE: WEIGHTED TRANSDUCTION

There may be scenarios where it is beneficial to more carefully choose which points to transduce rather than comparing every training point to every other training point. To this end, we outline a more sophisticated proposal, **weighted transduction**, that achieves this.

Sampling training pairs uniformly may violate the low rank structure needed for extrapolation discussed in Section 3.2. Consider a dataset of point clouds of various 3D objects in various orientations where the task is to predict their 3D grasping point. A predictor required to make predictions from similarities between point clouds of any two different objects and orientations may need to be far more complex than simpler predictors trained on only similar objects. Therefore, it may be useful to identify which training points to transduce, rather than transduce between all pairs of points during training.

Hence, we introduce **weighted transduction** described in Algorithm 2, which identifies which training points to transduce, rather than transduce between all pairs of points during training. During training, a transductive predictor \bar{h}_θ is trained together with a scalar bilinear *weighting function* ω_θ . \bar{h}_θ is trained to predict values as in the bilinear transduction algorithm, yet weighted by ω_θ . ω_θ predicts weights for data points and their potential anchors based on anchor point similarity and the anchor point, thereby deciding which points to transduce. We can either pre-train ω_θ , or jointly optimize ω_θ and \bar{h}_θ . At test time, we select an anchor point based on the learned weighting function, and make predictions via \bar{h}_θ . We show results on learning with weighted transduction on analytic functions in Appendix C.2 and on learning how to grasp from point clouds in Appendix C.3.

Algorithm 2 Weighted Transduction

- 1: **Input:** distance parameter ρ , training set $(x_1, y_1), \dots, (x_n, y_n)$, regularizer $\text{Reg}(\cdot)$
- 2: **Train:** Train θ on loss

$$\mathcal{L}(\theta) = \sum_{i=1}^n \sum_{j \neq i} \omega_\theta(x_i - x_j, x_j) \cdot \ell(\bar{h}_\theta(x_i - x_j, x_j), y_i) + \text{Reg}(\theta)$$

- 3: **Test:** for each new x_{test} , predict

$$y = \bar{h}_\theta(x_{\text{test}} - x_i, x_i), \text{ where } \mathbf{i} \sim \mathbb{P}[\mathbf{i} = i] \propto \omega_\theta(x_{\text{test}} - x_i, x_i)$$

C.2 ADDITIONAL RESULTS ON A SIMPLE 2-D EXAMPLE

Weighted transduction with a 2-D example We also considered a 2-D analytic function as shown in Fig 10. This function has a 2-D input (x_1, x_2) , and 2-D output (y_1, y_2) . The function is constructed by sampling a tiling of random values in the central block (between (1, 1) and (5, 5)). The remaining tiles are constructed by shifting the central tile by 6 units along the x_1 or x_2 direction. For tiles shifted in the x_1 -direction ((6, 0) or (0, 6)) the label has the same first dimension and a negated second dimension. For tiles shifted in the x_2 -direction ((6, 0) or (0, 6)) the label has the same second dimension and a negated first dimension. This has a particular symmetry in its construction where the problem has a very simple relationship between instances, but not every instance is related to every other instance (only those off by a constant offset). We show in Fig 10 that weighted transduction is particularly important in this domain as not every point corresponds to every other point. We perform weighted transduction (seeded by a few demonstration pairs), and this allows for extrapolation to OOS pairs. While this may seem like an OOC problem, we show through the standard bilinear (without reparameterization) comparison that there is no low rank structure between the dimensions but instead there is low rank structure on reparameterization. This shows the importance of using weighted transduction and the ability to find symmetries in the data with our proposed method.




Figure 10: Predictions for the 2-D analytic function from (x_1, x_2) to (y_1, y_2) . (**Top**) y_1 output values for inputs (x_1, x_2) , (**Bottom**) y_2 output values. (**Left to Right:**) In-distribution ground truth values, OOS ground truth values, neural network predictions, bilinear predictions, bilinear transduction predictions and weighted transduction predictions. Transduction weighting is important in this domain and it is able to discover problem symmetry. While this may seem like an OOC problem, since bilinear prediction directly doesn't work, the OOC view on this problem does not have low rank structure, while the OOS view does.

C.3 ADDITIONAL RESULTS ON SUPERVISED GRASP PREDICTION

Qualitative grasping prediction results For qualitative results of all models on the grasping point prediction task, please see Fig 11.

Bilinear transduction is comparable to an SE(3) equivariant architecture on OOS orientations and exceeds its performance on OOS scales We compare bilinear transduction to a method that builds in SE(3) equivariance domain knowledge into its architecture and operates on point clouds. Specifically, we compare to tensor field networks Thomas et al. (2018), and report the results in Table 2. To highlight that SE(3) equivariance architectures can extrapolate to new rotations but lack guarantees on scale, we train on 100 noisy samples of rotated mugs, bottles or teapots and on a separate set of scaled mugs, bottles or teapots by in-distribution values described in Appendix D.1. We test extrapolation to 50 mugs, bottles or teapots with out-of-sample orientations and scales sampled from the distribution described in Appendix D.1. We adapt the Tetris object classification tensor field network to learn grasp point prediction by training the mean of the last equivariant layer to predict a grasping point and applying the mean squared error (MSE) loss. Following the implementation in Thomas et al. (2018), we provide six manually labeled key points of the point cloud that uniquely identify the grasp point: two points on the handle and mug intersection and four points on the mug bottom. We compare to the linear, neural network and transduction baselines, as well as bilinear transduction. We fix these architectures to one layer and 32 units. Under various rotations, the tensor field network is able to extrapolate. For scaling however, the network error increases significantly, which is in accordance with the fact that scale transformation guarantees do not exist in SE(3) architectures. Bilinear transduction can extrapolate in both cases, is comparable to tensor field networks on OOS orientations and exceeds its performance on OOS scales.

Weighted grasp point prediction The grasp point prediction results for three objects described in Table 1 are trained with privileged knowledge of the object types at training time. I.e., the predictor at training time was provided only with pairs of same object categories but with different positions, orientations and scale. At test time an anchor was selected based on in-distribution similarities without knowledge of the test sample's object type. While this provides extra information, it is still not building in information about the distribution shift we would like to extrapolate to, i.e. the types of transformations the objects undergo: translation, rotation and scale. We apply the **weighted** transduction algorithm (Algorithm 2) to this domain and show that with more relaxed assumptions on the required object labels at training time we can learn weights and a predictor that perform as well as with privileged training (Table 3).

Table 2: Comparing bilinear transduction on grasp point prediction with tensor field networks [Thomas et al. \(2018\)](#), an **SE(3) equivariant** network for point clouds. We report mean and standard deviation for for 50 OOS samples.

Task	Tensor Field	Linear	Neural Net	Transduction	Ours
Mug rotation	0.009 ± 0.003	0.005 ± 0.001	0.066 ± 0.035	0.005 ± 0.001	0.01 ± 0.003
Mug scale	0.027 ± 0.012	0.006 ± 0.001	0.004 ± 0.001	0.004 ± 0.001	0.004 ± 0.001
Bottle rotation	0.008 ± 0.004	0.004 ± 0.004	0.085 ± 0.046	0.011 ± 0.026	0.004 ± 0.002
Bottle scale	0.019 ± 0.008	0.003 ± 0.001	0.003 ± 0.001	0.01 ± 0.006	0.003 ± 0.001
Teapot rotation	0.011 ± 0.006	0.003 ± 0.001	0.055 ± 0.03	0.017 ± 0.012	0.006 ± 0.002
Teapot scale	0.093 ± 0.039	0.003 ± 0.001	0.003 ± 0.001	0.003 ± 0.001	0.003 ± 0.001

For grasp point prediction, during training, we train a weighting function $\omega_\theta(x_j - x_i, x_i)$ to predict whether x_i should be transduced to x_j based on object-level labels of the data (transduce all mugs to other mugs and bottles to other bottles). Then we learn a predictor h_θ weighted by fixed ω_θ as described in Algorithm 2. At test time, given point x_j we select the training point x_i with highest value $\omega_\theta(x_j - x_i, x_i)$ and predict y_j via $h_\theta(x_j - x_i, x_i)$, without requiring privileged information about the object category at test time.




Figure 11: OOS grasping point prediction for bottles, mugs and teapots. Each **row** represents one test datapoint. Columns (**Left to Right**): linear, neural network, transductive and bilinear model predictions (red) and ground truth (green) grasp points. We show that baselines may predict infeasible grasping points (e.g. in the middle of an object) whereas bilinear transduction makes near accurate predictions.

C.4 ADDITIONAL RESULTS ON IMITATION LEARNING

Qualitative decision making results For qualitative results of all models on imitation learning domains, please view videos in the supplementary zip file.

Table 3: Grasping point prediction on three objects with various rotations, translations and scales with ground truth weighted training, standard training and **learned weighted** training for bilinear transduction with a fixed architecture. By training a weighting function with fewer labels, learned weighted bilinear transduction achieves similar performance to ground truth weighted training.

	Ground Truth Weighted	Bilinear Transduction	Learned Weighted
(Train)	0.007 ± 0.004	0.01 ± 0.005	0.006 ± 0.003
(OOS)	0.027 ± 0.014	0.068 ± 0.041	0.024 ± 0.014

C.5 CHOICE OF ARCHITECTURE

Periodic activation functions We provide further results on the grasp point prediction and imitation learning domains in Table 4. This experiment is similar to the one reported in Table 1, but with fourier pre-processing as the initial layer in each model. For the various domains, adding this pre-processing step is beneficial for some architecture-domain combinations, but does not surpass bilinear transduction on all domains for any method. Moreover, in most domains bilinear transduction achieves the best results, or close results to the best model. For implementation details see Section D.2.

Robustness across seeds and architectures: We show that our results are stable across the choice of architecture and seed. In Fig 12 we show the performance of our method and baselines over Meta-World tasks and architectures. In Table 5 we show the average performance of our method and baselines over the complex domains for fixed architecture parameters over three seeds.

Table 4: Mean and standard deviation over prediction (regression) or final state (sequential decision making) error for OOS samples and over a hyperparameter search with **fourier pre-processing**. While fourier activations are useful for some combinations of models and domains, it is not beneficial across all.

Task	Expert	Linear	Neural Net	DeepSets	Transduction	Ours
Mugs		0.119 ± 0.05	0.219 ± 0.114		0.3 ± 0.148	0.086 ± 0.081
Bottle		0.099 ± 0.041	0.139 ± 0.075		0.158 ± 0.096	0.019 ± 0.011
Teapot		0.112 ± 0.045	0.216 ± 0.107		0.301 ± 0.14	0.041 ± 0.026
All		0.142 ± 0.064	0.252 ± 0.119		0.33 ± 0.157	0.047 ± 0.033
Reach	0.006 ± 0.008	0.007 ± 0.006	0.075 ± 0.082	0.087 ± 0.092	0.117 ± 0.176	0.008 ± 0.008
Push	0.012 ± 0.001	0.182 ± 0.13	0.113 ± 0.124	0.121 ± 0.091	0.17 ± 0.12	0.014 ± 0.006
Slider	0.105 ± 0.066	0.38 ± 0.032	0.151 ± 0.081	0.241 ± 0.077	0.134 ± 0.114	0.256 ± 0.214
Adroit	0.035 ± 0.015	0.551 ± 0.508	0.398 ± 0.435	0.373 ± 0.195	0.341 ± 0.218	0.365 ± 0.31




Figure 12: Heatmap complementing Table 1 showing mean OOS errors on Meta-World. While some architectures can be suitable for some domains, our method is more robust to hyperparameter search and almost always achieves low error.

Table 5: We report mean and standard deviation error over evaluation samples and three **seeds**. For each domain, the first row is evaluated on in-distribution samples, the second row is on OOS samples. We show (1) bilinear transduction performs well in-distribution as well as OOS (2) our algorithm is stable across several seeds.

Task	Expert	Linear	Neural Net	DeepSets	Transduction	Ours
Grasping (Train)		0.06 ± 0.059	0.006 ± 0.004		0.006 ± 0.003	0.007 ± 0.003
Grasping OOS		0.141 ± 0.114	0.124 ± 0.074		0.099 ± 0.07	0.024 ± 0.014
Reach (Train)	0.006 ± 0.008	0.004 ± 0.005	0.005 ± 0.005	0.005 ± 0.005	0.004 ± 0.005	0.005 ± 0.005
Reach (OOS)		0.007 ± 0.006	0.013 ± 0.015	0.088 ± 0.075	0.007 ± 0.007	0.007 ± 0.006
Push (Train)	0.012 ± 0.001	0.22 ± 0.103	0.046 ± 0.049	0.01 ± 0.001	0.022 ± 0.011	0.016 ± 0.007
Push OOS		0.26 ± 0.063	0.355 ± 0.138	0.188 ± 0.124	0.337 ± 0.151	0.022 ± 0.016
Slider (Train)	0.105 ± 0.066	0.425 ± 0.077	0.344 ± 0.197	0.079 ± 0.057	0.32 ± 0.072	0.17 ± 0.092
Slider OOS		0.512 ± 0.06	0.438 ± 0.191	0.385 ± 0.288	0.55 ± 0.209	0.163 ± 0.126
Adroit (Train)	0.035 ± 0.015	0.66 ± 0.596	0.139 ± 0.231	0.103 ± 0.18	0.129 ± 0.215	0.047 ± 0.024
Adroit OOS		0.337 ± 0.075	0.703 ± 0.561	0.422 ± 0.241	0.773 ± 0.796	0.198 ± 0.255

D IMPLEMENTATION DETAILS

D.1 TRAIN/TEST DISTRIBUTIONS AND ANALYTIC FUNCTION DETAILS

Analytic functions: We used the following analytic functions for evaluation in Section 4.1.

Figure 4a. Periodic mixture of functions with different periods: We consider the period growing mixture of functions, and do a similar function but with two functions with different periods repeating. Let us define $x_v = x \bmod 9$, then the

$$f(x) = \begin{cases} x_m * \sin(10 * x_v) & \text{if } x_v < 1 \\ x_m * (\sin(10) + (x_v - 1)) & \text{if } 1 < x_v < 2 \\ x_m * (\sin(10) + 1 + (x_v - 2)^2) & \text{if } 2 < x_v < 3 \\ x_m * \sin(10 * \frac{(x_v - 3)}{2}) & \text{if } 3 < x_v < 5 \\ x_m * (\sin(10) + (\frac{x_v - 3}{2} - 1)) & \text{if } 5 < x_v < 7 \\ x_m * (\sin(10) + 1 + (\frac{x_v - 3}{2} - 2)^2) & \text{if } 7 < x_v < 9 \end{cases} \quad (\text{D.1})$$

We can see that this function mixes 2 different periods, to show that our framework can deal with varying functions as well.

Figure 4b. Sawtooth function: we use a classic sawtooth function

$$f(x) = \begin{cases} (x \bmod \text{Period}) * \frac{\text{Amplitude}}{\text{Period}} & \text{if } (\lfloor \frac{x}{\text{Period}} \rfloor) \bmod 2 == 0 \\ \text{Amplitude} - (x \bmod \text{Period}) * \frac{\text{Amplitude}}{\text{Period}} & \text{if } (\lfloor \frac{x}{\text{Period}} \rfloor) \bmod 2 == 1 \end{cases} \quad (\text{D.2})$$

Figure 4c. Randomly chosen polynomial: We sampled a random degree 8 polynomial: $f(x) = (x - 0.1)(x + 0.4)(x + 0.7)(x - 0.5)(x - 1.5)(x + 1.75)(x + 1)(x - 1.2)$

Figures 5 & 7. Periodic growing mixture of functions: Let us define $x_v = x \bmod 3$ and $x_m = \lfloor \frac{x}{3} \rfloor$

$$f(x) = \begin{cases} x_m * \sin(10 * x_v) & \text{if } x_v < 1 \\ x_m * (\sin(10) + (x_v - 1)) & \text{if } 1 < x_v < 2 \\ x_m * (\sin(10) + 1 + (x_v - 2)^2) & \text{if } 2 < x_v < 3 \end{cases} \quad (\text{D.3})$$

Figure 6. Shift equivariant mixture functions: Let us define $x_v = x \bmod 3$ and $x_m = \lfloor \frac{x}{3} \rfloor$

$$f(x) = \begin{cases} x_m * 3 + \sin(10 * x_v) & \text{if } x_v < 1 \\ x_m * 3 + (\sin(10) + (x_v - 1)) & \text{if } 1 < x_v < 2 \\ x_m * 3 + (\sin(10) + 1 + (x_v - 2)^2) & \text{if } 2 < x_v < 3 \end{cases} \quad (\text{D.4})$$

This represents a mixture of different functions - a sinusoid function, a linear function and a quadratic function that repeat and are equivariant on shifting. These are meant to show that bilinear transduction can capture equivariance.

Grasp point prediction: A training dataset is generated from rotating, translating and scaling a bottle, mug and teapot from the ShapeNet dataset [Chang et al. \(2015\)](#) with various parameters to predict a grasping point in \mathbb{R}^3 from their point clouds. A test dataset is constructed by applying the same transformations to the objects with OOS parameters. The point cloud in $\mathbb{R}^{N \times 3}$ is represented by a 12-tuple containing the point cloud’s mean point and its three PCA components in \mathbb{R}^3 . In [Fig 13](#) we plot the reduced point cloud state space representing point cloud mean and PCA components for in-distribution and OOS object orientations, positions and scales. We add Gaussian noise to the point cloud sampled from $\mathcal{N}(0, 0.005)$ and to the grasp point label sampled from $\mathcal{N}(0, 0.002)$. During training, objects are rotated around the z axis $\theta \in [0, 1.2\pi]$, translated by $(x, y) \in [0, 0.5] \times [0, 0.5]$ and scaled by a factor of $\alpha \in [0.7, 1.3]$. At test time these transformation parameters are sampled from the following ranges: $\theta \in [1.2\pi, 2\pi]$, $(x, y) \in [0.5, 0.7] \times [0.5, 0.7]$ and $\alpha \in [1.3, 1.6]$.

Meta-World: We evaluate on the reach-v2 and push-v2 tasks. For **reach-v2**, we reduce the observation space to a 6-tuple of the end effector 3D position and target 3D position. The action space is a 4-tuple of the end effector 3D position and the gripper’s degree of openness. During training, the 3D target position for reach-v2 is sampled from range $[0, 0.4] \times [0.7, 0.9] \times [0.05, 0.3]$ and at test time from range $[-0.4, 0] \times [0.7, 0.9] \times [0.05, 0.3]$. The gripper initial position is fixed to $[0, 0.6, 0.2]$. For **push-v2** the observation space is a 10-tuple of the end effector 3D position, gripper openness, object 3D position and target 3D position. The action space remains the same as in reach-v2. The training 3D target position is sampled from $[0, 0.3] \times [0.5, 0.7] \times \{0.01\}$ and at test time from $[-0.3, 0] \times [0.5, 0.7] \times \{0.01\}$. The gripper initial position is fixed to $[0, 0.4, 0.08]$ and the object initial position to $[0, 0.4, 0.02]$. For both environments, expert and evaluation trajectories are collected for 100 steps.

Adroit hand: We evaluate on the object relocation task in the ADROIT hand manipulation benchmark. The observation space is a 39-tuple of 30D hand joint positions, 3D object position, 3D palm position and 3D target position. The action space is a 30-tuple of 3D position and 3D orientation of the wrist, and 24D for finger positions commanding a PID controller. During training, the target location is sampled from $[0, 0.3] \times [0, 0.3] \times [0.15, 0.35]$ and at test time from $[-0.3, 0] \times [-0.3, 0] \times [0.15, 0.35]$. Expert and evaluation trajectories are collected for 200 steps.

Slider: We introduce an environment for controlling a brush to slide an object to a target. The observation space is an 18-tuple composed of a 9-tuple representing the brush 2D position, 3D object position and 4D quaternion, an 8-tuple representing the 2D brush velocity, 3D object velocity and 3D rotational object velocity, and the object mass. The action space is torque applied to the brush. During training, the object mass is sampled from $[60, 130]$ and at test time from $[5, 15]$. Expert and evaluation trajectories are collected for 200 steps. In [Figure 14](#) we plot the actions as a function of time steps for demonstrations, baselines described in [Section 4](#) and Bilinear Transduction for a set of fixed architecture parameters.

D.2 TRAINING DETAILS

D.2.1 DETAILS OF BILINEAR TRANSDUCTION IMPLEMENTATION

We describe the **bilinear transduction** algorithm implementation in detail. For grasp point prediction, during training we uniformly select an object $o \in \{\text{bottle, mug, teapot}\}$, and two instances of those objects, o_i, o_j that were rotated, translated and scaled with different parameters. We learn to transduce o_i to o_j , i.e. learn $h_\theta(o_j - o_i, o_i)$ to predict y_j , the grasping point of o_j . At test time, given test point x_j with an unknown object category, we select a training point x_i with an “in-distribution” difference $x_j - x_i$, as described in [Algorithm 1](#). We select ρ to be the closest difference $x_j - x_i$ to an in-distribution difference over all training points x_i . The differences generated by the optimization process are sampled uniformly to generate a smaller subset for comparison.

For the goal-conditioned imitation learning setting, during training we uniformly sample trajectories $\tau_i = \{(s_t^i, a_t^i)\}_{t=1}^T$ for horizon T specified in [Appendix D.1](#) and $\tau_j = \{(s_t^j, a_t^j)\}_{t=1}^T$. We further uniformly sample a time step t from $[0, T]$. We learn to transduce state $s_t^i \in \tau_i$ to $s_t^j \in \tau_j$. I.e., learn $h_\theta(s_t^j - s_t^i, s_t^i)$ to predict action a_t^j . At test time, we select an anchor trajectory based on the goal (Meta-World, Adroit) or mass (Slider) g_j . This is done by selecting the training trajectory τ_i

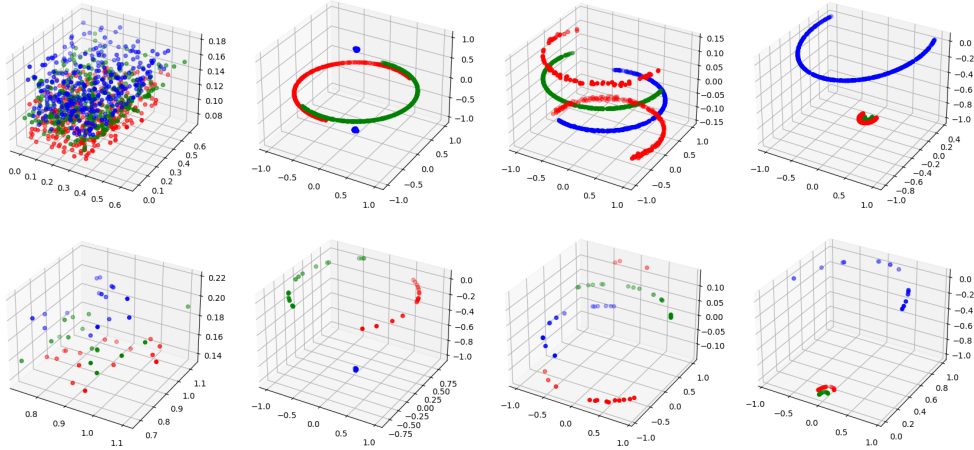


Figure 13: Reduced point cloud feature space in \mathbb{R}^{12} . From left to right: mean and three PCA components. Top: in distribution, bottom: OOS. Blue: bottles, red: mugs, green: teapots.

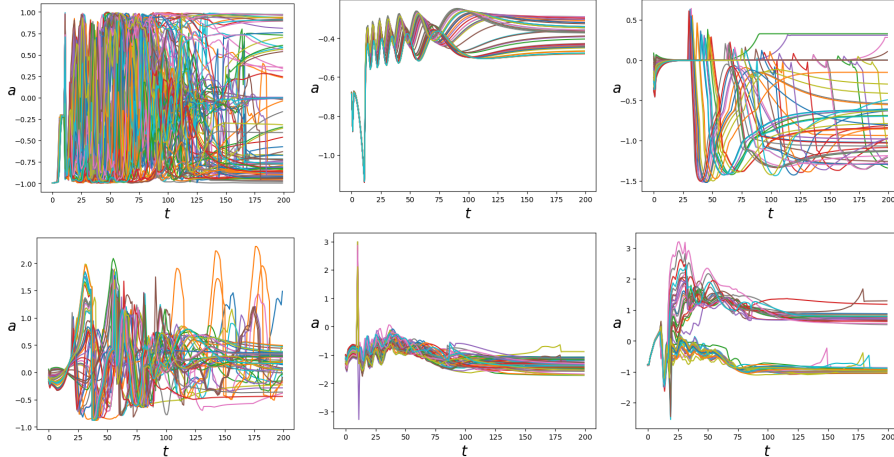


Figure 14: Slider torque actions as a function of time steps. Each color represents a different trajectory and mass. From left to right, top to bottom: in-support demonstrations and Linear, Neural Net, DeepSets, Transduction and Bilinear Transduction policies for OOS masses.

with goal or mass g_i that generates an “in-distribution” difference $g_j - g_i$ with the test point, as described in Algorithm 1. We select ρ to be the 10-th percentile of the in-distribution differences. The differences generated by the optimization process are approximated by generating differences from the training data. We then transduce each *state* in τ_i to predict actions for the test goal or mass. Given test state s_t^j , we predict a_t^j with $h_\theta(s_t^j - s_t^i, s_t^i)$, execute action a_t^j in the environment and observe state s_{t+1}^j . We complete this process for T steps to obtain τ_j .

D.2.2 DETAILS OF TRAINING DISTRIBUTIONS

We train on 1000 samples for all versions of the grasp point prediction domain: single objects, three objects and weighted transduction. In the sequential decision domains we train on N demonstrations, sequences of state-action pairs, $\{(x_i, y_i)\}_{i=1}^T$ where horizon T for each domain is specified in D.1. For Meta-World reach and push, $N = 1000$. For Slider and Adroit $N = 100$. All domains were evaluated on 50 in-distribution out-of-sample points and 50 OOS points. For each domain, all methods were trained on the same expert dataset and evaluated on the same in-distribution out-of-sample and OOS sets.

We collect expert demonstrations as follows. For grasp point prediction, we use the Neural Descriptor Fields (Simeonov et al., 2022) simulator to generate a feasible grasping point by a Franka Panda arm for three ShapeNet objects (bottle 2d1aa4e124c0dd5bf937e5c6aa3f9597, mug 61c10dccfa8e508e2d66cbf6a91063 and Utah teapot wss.ce04f39420c7c3e82fb82d326efadfe3). In Meta-World, we use the expert policies provided in the environment. For Adroit and Slider we train an expert policy using standard Reinforcement Learning (RL) methods - TRPO (Schulman et al., 2015) for Adroit and SAC (Haarnoja et al., 2018)

The weighting grasp point prediction network was provided with with 300 labels of positive pairs to be transduced and 300 negative pairs labeled with binary labels. This is a slightly more relaxed condition than labeling 1000 objects as in the standard version.

D.2.3 DETAILS OF MODEL ARCHITECTURES

For analytic domains, we use MLPs (both for NN and bilinear embeddings) with 3 layers of 1000 hidden units each with ReLU activations. We train all models with periodic activations of fourier features as described in Tancik et al. (2020).

Linear model: data points x are processed through a fully connected layer to output predictions y . **Neural Network:** data points x are processed through an mlp with k layers, n units each and ReLU activations for hidden layers to output predictions y . **DeepSets:** we refer to the elements of the state space excluding target position (Meta-World, Adroit) or mass (Slider) as observations. Observations and target position or mass are processed separately by two mlps with k layers, n units each and ReLU activations for hidden layers. Their embeddings are summed and processed by an mlp with one hidden layer with n units and ReLU activations to produce predictions y . **Transduction:** training point x_j and $x_i - x_j$ (for training or test point x_i) are concatenated and processed through an mlp with k layers, n units each and ReLU activations for hidden layers, to predict y_i . **Our architecture** **Bilinear Transduction:** training point x_j and $\Delta x_{ij} = x_i - x_j$ (for training or test point x_i) are embedded separately to a feature space in $\mathbb{R}^{d \cdot m}$, where d is the dimension of y_i , by two mlps with k layers, n units each and ReLU activations for hidden layers. Each predicted element of y_i is the dot product of m -sized segments of these embeddings as described in Eq. (3.5).

Further, we evaluate for all models with **Fourier Embedding** as a pre-processing step. We process inputs x , target or mass or Δx through a linear model which outputs $\times 40$ the number of inputs. We then multiply values by π and process through the sin function. This layer is optimized as the first layer of the model.

In Table 1 we search over number of layers $k \in \{2, 3\}$, and unit size $n \in \{32, 512, 1024\}$. The bilinear transduction embedding size is $m = 32$. In Tables 3 and 5 we set $k = 2$, $n = 32$ and $m = 32$.

The weighting function used for grasp point prediction is the bilinear transduction architecture with $k = 2$, $n = 128$ and $m = 32$.

D.2.4 DETAILS OF MODEL OPTIMIZATION

We train the analytic functions for 500 epochs, batch size 32, and Adam optimizer with learning rate $1e - 4$. We optimize the mean squared error (MSE) loss for regression.

We trained all more complex models for 5k epochs, batch size 32, with Adam (Kingma & Ba, 2014) optimizer and learning rate $1e-4$. We optimize the L2-norm loss function comparing between ground truth and predicted grasping points or actions for the sequential decision making domains. Each configuration of hyperparameters was ran and tested on one seed. We demonstrate the stability of our method across three seeds for a fixed set of hyperparameters in Table 5.

We train the weighted grasp point prediction for 5k epochs, batch size 16, and Adam optimizer with learning rate $1e - 4$. We optimize the MSE loss between the output and ground truth binary label indicating if a training point should be transduced to another training point. The weighting function did not require further finetuning jointly with the bilinear predictor.