

---

# Linear unit-tests for invariance discovery

---

**Benjamin Aubin**  
Facebook AI Research  
Paris, France

**Agnieszka Slowik**  
Facebook AI Research  
London, UK

**Martin Arjovsky**  
INRIA - PSL Research University  
Paris, France

**Leon Bottou**  
Facebook AI Research,  
New York, NY, 10003, USA

**David Lopez-Paz**  
Facebook AI Research  
Paris, France  
dlp@fb.com

## Abstract

There is an increasing interest in algorithms to learn invariant correlations across training environments. A big share of the current proposals find theoretical support in the causality literature but, how useful are they in practice? The purpose of this note is to propose six linear low-dimensional problems —“unit tests”— to evaluate different types of out-of-distribution generalization in a precise manner. Following initial experiments, none of three recently proposed alternatives passes all tests. By providing the code to automatically replicate all the results in this manuscript (<https://www.github.com/facebookresearch/InvarianceUnitTests>), we hope that our unit tests become a standard stepping stone for researchers in out-of-distribution generalization.

## 1 Introduction

Machine learning systems crumble when deployed in conditions different to those of training [Szegedy et al., 2013, Rosenfeld et al., 2018, Alcorn et al., 2019]. To address this issue, recent works in causality [Peters et al., 2015, Arjovsky et al., 2019, Parascandolo et al., 2020] propose to learn correlations invariant across multiple training distributions, and to use those correlations as a proxy for out-of-distribution generalization. However, as we will show, these algorithms perform poorly across a catalog of simple low-dimensional linear problems. Therefore, our contribution is a standardized set of six “unit tests” that researchers can bear in mind when proposing new solutions for out-of-distribution generalization.

Causal learning algorithms such as Invariant Causal Prediction [Peters et al., 2015, ICP] and Invariant Risk Minimization [Arjovsky et al., 2019, IRM] consume several training datasets, each of them possibly produced by the same “structural equation model” operating under a different “valid interventions”. A structural equation model is a list of equations describing how variables influence each other to take their values [Pearl, 2009, Peters et al., 2017]. An intervention perturbs one or more of these equations, and it is “valid” as long as it does not modify the conditional expectation of the target variable given its direct causal parents [Arjovsky et al., 2019]. Then, if the interventions producing our training datasets are diverse, invariant correlations should pertain to the fixed causal mechanism of the target variable. This suggests the possibility of learning about the causal structure of data by searching for statistical invariances. In turn, these invariances enable out-of-distribution generalization: if the causal mechanism of the target variable is invariant across diverse training conditions, we may rely on them to perform robustly at novel test conditions.

Unfortunately, invariances are often more difficult to capture than spurious correlations. As an example, consider the task of classifying pictures of cows and camels [Arjovsky et al., 2019]. A “green detector” solves this task to a great extent, since almost all pictures of cows contain green pastures, and almost all pictures of camels show beige sandy landscapes. Following the principle of least effort [Geirhos et al., 2020], learning machines cling to the textural “green-cow” spurious correlation, rather than discovering the shapes that make a cow a cow. Predictors absorbing these “distractor”, “shortcut”, or “bait” correlations fail when deployed under novel conditions.

Although designed to address this very issue, causal learning algorithms often fail to capture causal invariances in data. The purpose of this note is to share six linear problems that illustrate this phenomena. Each of these problems contains an invariant causal correlation (*inv*) that we would like to learn, as well as a spurious correlation (*spu*) that we would like to discard. In every dataset, empirical risk minimization absorbs the spurious correlation from the training data, failing to generalize to novel conditions. By releasing the code to automatically replicate all the results in this manuscript, we hope that these “unit tests” become a standard guide when developing new out-of-distribution generalization algorithms.

## 2 Problems

For each problem, we collect datasets  $D_e = \{(x_i^e, y_i^e)\}_{i=1}^{n_e}$  containing  $n_e$  samples for  $n_{\text{env}}$  environments  $e \in \mathcal{E} = \{E_j\}_{j=1}^{n_{\text{env}}}$ . The input feature vector  $x^e = (x_{\text{inv}}^e, x_{\text{spu}}^e) \in \mathbb{R}^d$  contains features  $x_{\text{inv}}^e \in \mathbb{R}^{d_{\text{inv}}}$  that elicit invariant correlations, as well as features  $x_{\text{spu}}^e \in \mathbb{R}^{d_{\text{spu}}}$  that elicit spurious correlations. Our goal is to construct invariant predictors that estimate the target variable  $y^e$  by relying on  $x_{\text{inv}}^e$ , while ignoring  $x_{\text{spu}}^e$ . To measure the extent to which an algorithm ignores the features  $x_{\text{spu}}^e$ , we sample a *train* split, a *validation* split, and a *test* split per problem and environment. Both train and validation splits are built by sampling the structural equations outlined below for each problem and environment. The test split is built analogously, but the features  $x_{\text{spu}}^e$  are shuffled at random across examples. This way, only those predictors ignoring the shortcut correlations provided by  $x_{\text{spu}}^e$  will achieve minimal test error.

Finally, let  $\mathcal{N}_d(\mu, \sigma^2)$  be the  $d$ -dimensional Gaussian distribution with mean vector  $\mu$  and diagonal covariance matrix with diagonal elements  $\sigma^2$ . We fix  $n_e = 10^4$  and we mainly illustrate our results in the default case  $n_{\text{env}} = 3$  with  $(d, d_{\text{inv}}, d_{\text{spu}}) = (10, 5, 5)$ .

### 2.1 Example1: regression from causes and effects

A linear least-squares regression problem where features contain causes and effects of the target variable, as proposed by Arjovsky et al. [2019]. By virtue of considering only valid interventions, the mapping from the causes to the target variable is invariant across environments. Conversely, the mapping from the target variable to the effects changes across environments. To construct the datasets  $D_e$  for every  $e \in \mathcal{E}$  and  $i = 1, \dots, n_e$ , sample:

$$\begin{aligned} x_{\text{inv},i}^e &\sim \mathcal{N}_{d_{\text{inv}}}(0, (\sigma^e)^2), & x_i^e &\leftarrow (x_{\text{inv},i}^e, x_{\text{spu},i}^e), \\ \tilde{y}_i^e &\sim \mathcal{N}_{d_{\text{inv}}}(W_{yx} x_{\text{inv},i}^e, (\sigma^e)^2), & y_i^e &\leftarrow \frac{2}{d} \cdot \mathbf{1}_{d_{\text{inv}}}^\top \tilde{y}_i^e; \\ x_{\text{spu},i}^e &\sim \mathcal{N}_{d_{\text{spu}}}(W_{xy} \tilde{y}_i^e, 1), \end{aligned}$$

where the matrices  $W_{yx} \in \mathbb{R}^{d_{\text{inv}} \times d_{\text{inv}}}$ ,  $W_{xy} \in \mathbb{R}^{d_{\text{spu}} \times d_{\text{inv}}}$  are drawn i.i.d from the Gaussian normal distribution. The first environment variables are fixed to  $(\sigma^{e=E_0})^2 = 0.1$ ,  $(\sigma^{e=E_1})^2 = 1.5$ , and  $(\sigma^{e=E_2})^2 = 2$ . For  $n_{\text{env}} > 3$  and  $j \in [3 : n_{\text{env}} - 1]$ , the extra environments  $(\sigma^{e=E_j})^2$  are drawn uniformly from  $\text{Unif}(10^{-2}, 10)$ .

**Challenges** First, the distribution of the feature  $x_{\text{inv}}^e$  eliciting the invariant correlation changes across environments. This disallows the use of domain-adversarial methods [Ganin et al., 2016], which seek features with matching distribution across environments. Second, the distribution of the residuals varies across environments, disallowing the use of ICP [Peters et al., 2015]. Third, since the target variable is continuous, conditional domain-adversarial techniques [Li et al., 2018] do not apply easily. This example is solved by IRM, and it is analogous to a linear regression variant of the ColorMNIST task [Arjovsky et al., 2019].

## 2.2 Example2: cows versus camels

In the spirit of [Beery et al., 2018, Arjovsky et al., 2019], we add a binary classification problem to imitate the introductory example “most cows appear in grass and most camels appear in sand”. Let:

$$\begin{aligned} \mu_{\text{cow}} &\sim 1_{d_{\text{inv}}}, & \mu_{\text{camel}} &= -\mu_{\text{cow}}, & \nu_{\text{animal}} &= 10^{-2}, \\ \mu_{\text{grass}} &\sim 1_{d_{\text{spu}}}, & \mu_{\text{sand}} &= -\mu_{\text{grass}}, & \nu_{\text{background}} &= 1. \end{aligned}$$

To construct the datasets  $D_e$  for every  $e \in \mathcal{E}$  and  $i = 1, \dots, n_e$ , sample:

$$j_i^e \sim \text{Categorical}(p^e s^e, (1-p^e)s^e, p^e(1-s^e), (1-p^e)(1-s^e));$$

$$\begin{aligned} x_{\text{inv},i}^e &\sim \begin{cases} \mathcal{N}_{d_{\text{inv}}}(0, 10^{-1}) + \mu_{\text{cow}} \cdot \nu_{\text{animal}} & \text{if } j_i^e \in \{1, 2\}, \\ \mathcal{N}_{d_{\text{inv}}}(0, 10^{-1}) + \mu_{\text{camel}} \cdot \nu_{\text{animal}} & \text{if } j_i^e \in \{3, 4\}, \end{cases} & x_i^e &\leftarrow (x_{\text{inv},i}^e, x_{\text{spu},i}^e); \\ x_{\text{spu},i}^e &\sim \begin{cases} \mathcal{N}_{d_{\text{spu}}}(0, 10^{-1}) + \mu_{\text{grass}} \cdot \nu_{\text{background}} & \text{if } j_i^e \in \{1, 4\}, \\ \mathcal{N}_{d_{\text{spu}}}(0, 10^{-1}) + \mu_{\text{sand}} \cdot \nu_{\text{background}} & \text{if } j_i^e \in \{2, 3\}, \end{cases} & y_i^e &\leftarrow \begin{cases} 1 & \text{if } \mathbf{1}_{d_{\text{inv}}}^\top x_{\text{inv},i}^e > 0, \\ 0 & \text{else;} \end{cases} \end{aligned}$$

where the environment foreground/background probabilities are  $p^{e=E_0} = 0.95$ ,  $p^{e=E_1} = 0.97$ ,  $p^{e=E_2} = 0.99$  and the cow/camel probabilities are  $s^{e=E_0} = 0.3$ ,  $s^{e=E_1} = 0.5$ ,  $s^{e=E_2} = 0.7$ . For  $n_{\text{env}} > 3$  and  $j \in [3 : n_{\text{env}} - 1]$ , the extra environment variables are respectively drawn according to  $p^{e=E_j} \sim \text{Unif}(0.9, 1)$  and  $s^{e=E_j} \sim \text{Unif}(0.3, 0.7)$ .

**Challenges** Achieving zero *population* error while using only  $x_{\text{inv}}^e$  requires learning large weights. This is difficult when using gradient descent methods, or most forms of regularization. As the dimension of the feature space grows, the probability of achieving zero *training* error using only  $x_{\text{spu}}^e$  increases rapidly. This means that invariance penalties based on training error, such as IRM, may accept solutions using spurious features.

## 2.3 Example3: small invariant margin

A linear version of the spiral binary classification problem proposed by [Parascandolo et al., 2020], where the first two dimensions offer an invariant, small-margin linear decision boundary. The rest of the dimensions offer a changing, large-margin linear decision boundary. Let  $\gamma = 0.1 \cdot 1_{d_{\text{inv}}}$ , and  $\mu^e \sim \mathcal{N}_{d_{\text{spu}}}(0, 1)$ , for all environments.

To construct the datasets  $D_e$  for every  $e \in \mathcal{E}$  and  $i = 1, \dots, n_e$ , sample:

$$\begin{aligned} y_i^e &\sim \text{Bernoulli}\left(\frac{1}{2}\right), \\ x_{\text{inv},i}^e &\sim \begin{cases} \mathcal{N}_{d_{\text{inv}}}(+\gamma, 10^{-1}) & \text{if } y_i^e = 0, \\ \mathcal{N}_{d_{\text{inv}}}(-\gamma, 10^{-1}) & \text{if } y_i^e = 1; \end{cases} & x_i^e &\leftarrow (x_{\text{inv},i}^e, x_{\text{spu},i}^e). \\ x_{\text{spu},i}^e &\sim \begin{cases} \mathcal{N}_{d_{\text{spu}}}(+\mu^e, 10^{-1}) & \text{if } y_i^e = 0, \\ \mathcal{N}_{d_{\text{spu}}}(-\mu^e, 10^{-1}) & \text{if } y_i^e = 1; \end{cases} \end{aligned}$$

**Challenges** We can solve this problem to zero *population* error with high probability using  $x_{\text{spu}}^e$  alone. Things complicate further, as solving this task using  $x_{\text{inv}}^e$  forcefully incurs a small amount of *population* error. Therefore, learning algorithms should learn to sacrifice training error to realize that  $x_{\text{inv}}^e$  lead to the same *maximum margin* classifier across environments (even though the varying margin based on  $x_{\text{spu}}^e$  is larger!). While the predictor based on  $x_{\text{inv}}^e$  is the optimal in terms of worst-case out-of-distribution generalization, it is not a causal predictor of the target variable.

## 2.4 Scrambled variations

We define three additional problems: `example1s`, `example2s`, and `example3s`. These are “scrambled” variations of the three problems described above, respectively. Scrambled variations build observed datasets  $D^e = \{(S^\top x_i^e, y_i^e)\}_{i=1}^{n_e}$ , where  $S \in \mathbb{R}^{d \times d}$  is a random rotation matrix fixed for all environments  $e \in \mathcal{E}$ . Observing a scrambled version of the variables appearing in the structural equation models requires algorithms to learn a (linear) feature representation under which the desired invariance should be elicited. Because of this reason, we do not compare to brute-force feature selection methods, such as ICP [Peters et al., 2015].

Table 1: Test errors for all algorithms, datasets, and environments for  $(d_{\text{inv}}, d_{\text{spu}}, n_{\text{env}}) = (5, 5, 3)$ . Example 1 and Example 1s errors are in MSE; all others are classification errors.

	ANDMask	ERM	IGA	IRMv1	Oracle
Example1.E0	$0.11 \pm 0.04$	$1.62 \pm 0.60$	$4.47 \pm 1.16$	$0.20 \pm 0.04$	$0.05 \pm 0.00$
Example1.E1	$11.39 \pm 0.18$	$14.25 \pm 1.52$	$18.46 \pm 2.14$	$11.98 \pm 0.75$	$11.27 \pm 0.17$
Example1.E2	$20.28 \pm 0.30$	$24.22 \pm 2.34$	$29.48 \pm 3.19$	$21.27 \pm 1.34$	$19.93 \pm 0.31$
Example1s.E0	$0.07 \pm 0.01$	$1.61 \pm 0.59$	$4.55 \pm 1.79$	$0.19 \pm 0.04$	$0.05 \pm 0.00$
Example1s.E1	$12.13 \pm 0.80$	$14.23 \pm 1.49$	$18.68 \pm 3.37$	$11.92 \pm 0.69$	$11.24 \pm 0.19$
Example1s.E2	$21.52 \pm 1.42$	$24.14 \pm 2.39$	$29.81 \pm 4.78$	$21.08 \pm 1.31$	$20.06 \pm 0.37$
Example2.E0	$0.42 \pm 0.02$	$0.40 \pm 0.01$	$0.43 \pm 0.00$	$0.43 \pm 0.00$	$0.00 \pm 0.00$
Example2.E1	$0.49 \pm 0.03$	$0.47 \pm 0.01$	$0.50 \pm 0.00$	$0.50 \pm 0.00$	$0.00 \pm 0.00$
Example2.E2	$0.42 \pm 0.02$	$0.40 \pm 0.01$	$0.42 \pm 0.01$	$0.42 \pm 0.01$	$0.00 \pm 0.00$
Example2s.E0	$0.43 \pm 0.01$	$0.43 \pm 0.01$	$0.43 \pm 0.01$	$0.43 \pm 0.01$	$0.00 \pm 0.00$
Example2s.E1	$0.50 \pm 0.00$	$0.50 \pm 0.00$	$0.50 \pm 0.00$	$0.50 \pm 0.00$	$0.00 \pm 0.00$
Example2s.E2	$0.42 \pm 0.01$	$0.42 \pm 0.01$	$0.42 \pm 0.01$	$0.42 \pm 0.01$	$0.00 \pm 0.00$
Example3.E0	$0.35 \pm 0.22$	$0.48 \pm 0.09$	$0.47 \pm 0.10$	$0.49 \pm 0.07$	$0.00 \pm 0.00$
Example3.E1	$0.36 \pm 0.22$	$0.48 \pm 0.07$	$0.48 \pm 0.08$	$0.49 \pm 0.06$	$0.00 \pm 0.00$
Example3.E2	$0.32 \pm 0.22$	$0.47 \pm 0.12$	$0.46 \pm 0.12$	$0.48 \pm 0.07$	$0.00 \pm 0.00$
Example3s.E0	$0.45 \pm 0.13$	$0.48 \pm 0.08$	$0.48 \pm 0.09$	$0.49 \pm 0.07$	$0.00 \pm 0.00$
Example3s.E1	$0.49 \pm 0.05$	$0.49 \pm 0.05$	$0.48 \pm 0.07$	$0.49 \pm 0.06$	$0.00 \pm 0.00$
Example3s.E2	$0.46 \pm 0.12$	$0.47 \pm 0.09$	$0.47 \pm 0.11$	$0.48 \pm 0.07$	$0.00 \pm 0.00$




Figure 1: Test error averaged across environments (E0, E1, E2) for  $(d_{\text{inv}}, d_{\text{spu}}, n_{\text{env}}) = (5, 5, 3)$ .

### 3 Experiments

We provide an initial set of experiments evaluating the following algorithms on our six problems: Empirical Risk Minimization [Vapnik, 1998, ERM] minimizes the error on the union of all the training splits. Invariant Risk Minimization [Arjovsky et al., 2019, IRMv1] finds a representation of the features such that the optimal classifier, on top of that representation, is the identity function for all environments. Inter-environmental Gradient Alignment [Koyama and Yamaguchi, 2020, IGA] minimizes the error on the training splits while reducing the variance of the gradient of the loss per environment. AND-mask [Parascandolo et al., 2020] minimizes the error on the training splits by updating the model on those directions where the sign of the gradient of the loss is the same for most environments. Oracle is a version of ERM where all data splits contain randomized  $x_{\text{spu}}^e$ , and therefore are trivial to ignore. The purpose of this method is to understand the achievable upper bound performance in our problems.

For each algorithm, we run a random hyper-parameter search of 20 trials. We trained each algorithm and hyper-parameter trial on the train splits of all environments, for  $10^4$  full-batch Adam [Kingma and Ba, 2015] updates. We choose the hyper-parameters trial that minimizes the error on the validation splits of all environments. Finally, we report the error of these selected models on the test splits. To provide error bars, this entire process, including data sampling, is repeated 50 times. We refer the reader to our code to learn about the hyper-parameter search distributions for each algorithm.

#### 3.1 Default results

Table 1 shows that no method is able to achieve a performance close to the Oracle on any of the proposed problems. The only exception are IRMv1 and ANDMask on example1. This illustrates that current causal learning algorithms are unable to capture invariances, even in low-dimensional linear problems. The results averaged over the different environments are plotted in Fig. 1.




Figure 2: Test error averaged across environments for ANDMask, ERM, IGA, IRMv1 and Oracle on the unit-tests as **(top)** a function of the ratio  $\delta_{env} = n_{env}/d_{spu}$  at fixed dimensions  $(d_{inv}, d_{spu}) = (5, 5)$ ; and as **(bottom)** a function of  $\delta_{spu} = d_{spu}/d_{inv}$  for  $(d_{inv}, n_{env}) = (5, 3)$ .

### 3.2 Varying the number of environments

What is the role of the number of environments  $n_{env}$  on generalization? We define the ratio  $\delta_{env} = \frac{n_{env}}{d_{spu}}$  between the number of environments and the number of spurious dimensions. We run experiments with the same procedure described above for  $n_{env} \in [2 : 10]$ , and a fixed number of spurious dimensions  $d_{spu} = d_{inv} = 5$ . Figure 2 **(top)** show average test errors for all algorithms. Notably, IGA performs no better than ERM, except on Example2 where it improves drastically when increasing  $\delta_{env}$ ; however, this increase bears no effect on the scrambled version Example2s. On Example1 and Example1s, both ANDMask and IRMv1 approach closely Oracle’s performances, while on Example2 and Example2s simple ERM outperforms them. On the contrary, ANDMask and IRMv1 achieve good performances on Example3. They approach optimality for  $n_{env} \simeq d_{spu} + 1$ , since in this case no invariant boundary can solve the problem using  $x_{spu}$  alone. IRMv1 performances do not suffer due to scrambling, while ANDMask collapses on Example3s.

### 3.3 Varying the number of spurious dimensions

We perform another ablation by fixing the number of environments  $n_{env} = 3$  and the number of invariant dimensions  $d_{inv} = 5$ , while varying the number of spurious dimensions  $\delta_{spu} = \frac{d_{spu}}{d_{inv}}$ . We observe that for Example1 and Example1s, ANDMask and IRMv1 do not suffer when adding spurious dimensions, while IGA crumbles as soon as a single spurious feature is added. As expected, on Example3(s) and Example3s, increasing the number of spurious dimensions while keeping the number of environments fixed decreases the performance of all algorithms. Example2 and Example2s show the same phenomena for  $\delta_{spu} \leq 1$ .

## 4 Outlook

We propose a battery of “unit-tests” to surgically evaluate different types of out-of-distribution generalization abilities of machine learning algorithms. While admittedly synthetic, our collection of problems attempts to cover a wide range of challenging distributional discrepancies that may arise across training and testing conditions. We invite researchers to use and extend this set of problems to learn about the strengths and shortcomings of new algorithms in a transparent and standardized manner.

## References

- Michael A Alcorn, Qi Li, Zhitao Gong, Chengfei Wang, Long Mai, Wei-Shinn Ku, and Anh Nguyen. Strike (with) a pose: Neural networks are easily fooled by strange poses of familiar objects. *CVPR*, 2019.
- Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.
- Sara Beery, Grant Van Horn, and Pietro Perona. Recognition in terra incognita. *ECCV*, 2018.
- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *JMLR*, 2016.
- Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. Shortcut learning in deep neural networks. *arXiv preprint arXiv:2004.07780*, 2020.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *ICLR*, 2015.
- Masanori Koyama and Shoichiro Yamaguchi. Out-of-distribution generalization with maximal invariant predictor. *arXiv preprint arXiv:2008.01883*, 2020.
- Ya Li, Xinmei Tian, Mingming Gong, Yajing Liu, Tongliang Liu, Kun Zhang, and Dacheng Tao. Deep domain generalization via conditional invariant adversarial networks. *ECCV*, 2018.
- Giambattista Parascandolo, Alexander Neitz, Antonio Orvieto, Luigi Gresele, and Bernhard Schölkopf. Learning explanations that are hard to vary. *arXiv preprint arXiv:2009.00329*, 2020.
- Judea Pearl. *Causality*. Cambridge university press, 2009.
- Jonas Peters, Peter Bühlmann, and Nicolai Meinshausen. Causal inference using invariant prediction: identification and confidence intervals. *arXiv preprint arXiv:1501.01332*, 2015.
- Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference*. 2017.
- Amir Rosenfeld, Richard Zemel, and John K Tsotsos. The elephant in the room. *arXiv preprint arXiv:1808.03305*, 2018.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- Vladimir Vapnik. *Statistical learning theory* wiley. *New York*, 1998.