

InferPy: Probabilistic Modeling Made Easy

Rafael Cabañas, Andrés R. Masegosa, Antonio Salmerón

University of Almería, ES-04120 Almería, Spain

Abstract

InferPy is a high-level Python API for probabilistic modeling built on top of Edward and Tensorflow. InferPy, which is strongly inspired by Keras, focuses on being user-friendly by using an intuitive set of abstractions that make easy to deal with complex probabilistic models. It should be seen as an interface rather than a standalone machine-learning framework. In general, InferPy has the focus on enabling flexible data processing, easy-to-code probabilistic modeling, scalable inference and robust model validation.

Keywords: keyword 1, keyword 2, keyword 3

1. Introduction

Machine learning (ML) [1] is a fundamental part of many artificial intelligence techniques [2], and the key of many innovative applications. Unfortunately, for a company or an institution, the development of ML models specific to their problems requires enormous efforts [3]. For this reason, probabilistic programming languages (PPLs) [4] are an active area of research. PPLs offer the same advantages to the ML community that high-level programming languages offered to software developers fifty years ago [5]. Programmers could specialize in model development while ML experts could focus their efforts on developing reusable inference engines. Thus, the number of non-experts who can create applications using a PPL could increase. Special attention requires those PPLs which exploit recent advances in probabilistic inference for defining probabilistic models containing deep neural networks [6, 7]. These PPLs rely on deep learning libraries like Tensorflow

Email addresses: `rcabanass@ual.es` (Rafael Cabañas), `andresmasegosa@ual.es` (Andrés R. Masegosa), `antonio.salmeron@ual.es` (Antonio Salmerón)

15 [8]. The main drawback of these approaches is the high complexity of the
 16 provided abstractions, specially those centered around the definition of prob-
 17 ability distributions over multidimensional Tensors.

18 InferPy¹ tries to address these issues by defining a user-friendly API which
 19 trades-off model complexity with ease of use. Complex operations over Ten-
 20 sor objects are hidden to the user. Similarly, Edward’s flexible approach to
 21 probabilistic inference demands to provide specific details such as the varia-
 22 tional family. Again, InferPy gives the possibility to hide all this information
 23 and make inference with a single line of code. As InferPy uses Tensorflow as
 24 computing engine, all the parallelization details are hidden to the user.

25 2. Background

26 InferPy focuses on *hierarchical probabilistic models* structured in two lay-
 27 ers: (i) a *prior model* defining a joint distribution $p(\mathbf{w})$ over the global pa-
 28 rameters of the model (\mathbf{w} can be a single random variable or a bunch of
 29 random variables with any given dependency structure); (ii) a *data or ob-*
 30 *servaion model* defining a joint conditional distribution $p(\mathbf{x}, \mathbf{z}|\mathbf{w})$ over the
 31 observed quantities \mathbf{x} , and the the local hidden variables \mathbf{z} governing the
 32 observation \mathbf{x} . As a running example, Figure 1 shows a model of this type.




Figure 1: Mixture of K d -dimensional Gaussian distributions learned from N observations.

33 3. Software Framework

34 3.1. Model Definition

35 In InferPy, models are specified using a simple language of random vari-
 36 ables, which are grouped in a *probabilistic model* object (i.e., defined using
 37 the construct `with inf.ProbModel() as m:`) defining a joint distribution
 38 over observable and hidden variables $p(\mathbf{w}, \mathbf{z}, \mathbf{x})$. As an example, we provide
 39 in Figure 2 how the model of Figure 1 would be defined in InferPy.

¹Home: inferpy.readthedocs.io; Source: github.com/PGM-Lab/InferPy

```

1 ## model definition ##
2 with inf.ProbModel() as model:
3
4     # prior distributions
5     with inf.replicate(size=K):
6         mu = inf.models.Normal(loc=0, scale=1, dim=d)
7         sigma = inf.models.InverseGamma(1, 1, dim=d)
8     p = inf.models.Dirichlet(np.ones(K)/K)
9
10    # define the generative model
11    with inf.replicate(size=N):
12        z = inf.models.Categorical(probs = p)
13        x = inf.models.Normal(mu[z], sigma[z], observed=True, dim=d)

```

Figure 2: InferPy code for the Mixture of Gaussians model of Figure 1.

40 InferPy allows to specify our model in a single sample-basis, resembling
41 the standard *plateau notation*, with the `with inf.replicate(size=N)` con-
42 struct (Line 5). The dimension N is the number of *replicas* of this part of
43 the model. The dimension of each variable can be specified either using the
44 input parameter `dim` (Line 6), or by the length of the distribution param-
45 eters (e.g., other InferPy variable, NumPy’s ndarray [9], a tensor or a Python
46 list). For example, variable `x` in the previous code contains N replicas of d
47 independent Gaussian distributions and, in consequence, has two dimensions
48 (i.e., $shape = [N, d]$).

49 Like in Edward, each random variable y is associated to a tensor y^* rep-
50 resenting a sample from its distribution. Note that when operating on y , the
51 operation is indeed done on y^* . In the previous code, the mean (i.e., `loc`) of
52 `x` is a sample from the distribution obtained by indexing `mu` with a sample
53 from `z`. Any variable defined in InferPy encapsulates an equivalent one in
54 Edward, which can be obtain by accessing the property `dist`. For simplicity,
55 the user does not deal with tensor objects unless it is explicitly specified, e.g.:
56 `z.sample()` returns an array of samples while `z.sample(tf_run=False)` al-
57 lows to obtain the equivalent (lazily evaluated) Tensor object.

58 3.2. Approximate Inference

59 InferPy directly relies on top of Edward’s inference engine. In particu-
60 lar, InferPy inherits Edward’s approach and considers approximate inference
61 solutions in which the task is to approximate the posterior with a simpler
62 distribution q . Unlike Edward, InferPy offers the possibility to hide all these

63 details about the definition of this q distribution, making the inference more
64 simple for non-advanced users. Figure 4 shows the code for making inference
65 in the model defined in the previous section.

```
1 # compile and fit the model with training data
2 data = {x: x_train}
3 model.compile(infMethod="MCMC")
4 model.fit(data)
5 # print the posterior
6 print(model.posterior(mu))
```

Figure 3: Code for making inference in the Mixture of Gaussian model of Figure 2.

66 4. Comparison with Edward

67 The analogous Edward code for defining and making inference in a mix-
68 ture of Gaussians, which can be found in our online documentation², has
69 some drawbacks compared to the code in InferPy (Figures 2 and 4). First, the
70 model definition is more complex because this is not done in a single-sample
71 basis. This can be specially problematic when defining the dependencies
72 among variables. For example, the mean of \mathbf{x} is specified using the func-
73 tion `tf.gather` which is not always intuitive, i.e. `loc=tf.gather(mu,z)`.
74 Secondly, Edward requires to have a strong knowledge about the inference
75 algorithms for specifying all its parameters. For the running example, a q
76 and g variable is defined for each latent variable in the model. For variable
77 μ , this is done as follows.

```
1 qmu = ed.models.Empirical(params=tf.get_variable("qmu/prm", [T,K,d],
2                                     initializer=tf.zeros_initializer()))
3 gmu = ed.models.Normal(loc=tf.ones([K,d]), scale=tf.ones([K,d]))
```

Figure 4: Edward’s code for defining the q distribution for the model of Figure 2.

78 5. Conclusions

79 We have briefly presented InferPy, a high-level API for probabilistic mod-
80 eling built on top of Edward and Tensorflow. The use of intuitive abstractions

²https://inferpy.readthedocs.io/en/latest/notes/inf_vs_ed.html

81 such as the *plateau notation* simplifies the task of defining complex hierarchical
82 probabilistic models. In the future, we aim to fully integrate InferPy with
83 Keras, allowing simple probabilistic modeling with deep neural networks.

84 **Acknowledgements**

85 Authors have been jointly supported by the Spanish Ministry of Science,
86 Innovation and Universities and by the FEDER under the projects TIN2015-
87 74368-JIN, and TIN2016-77902-C3-3-P.

- 88 [1] K. P. Murphy, Machine learning: A probabilistic perspective. adaptive
89 computation and machine learning (2012).
- 90 [2] S. J. Russell, P. Norvig, Artificial intelligence: a modern approach,
91 Malaysia; Pearson Education Limited,, 2016.
- 92 [3] Z. Ghahramani, Probabilistic machine learning and artificial intelligence,
93 Nature 521 (7553) (2015) 452.
- 94 [4] A. D. Gordon, T. A. Henzinger, A. V. Nori, S. K. Rajamani, Probabilistic
95 programming, in: Proceedings of the on Future of Software Engineering,
96 FOSE 2014, ACM, 2014, pp. 167–181.
- 97 [5] R. L. Wexelblat, History of programming languages, Academic Press,
98 2014.
- 99 [6] D. Tran, A. Kucukelbir, A. B. Dieng, M. Rudolph, D. Liang, D. M. Blei,
100 Edward: A library for probabilistic modeling, inference, and criticism,
101 arXiv preprint arXiv:1610.09787.
- 102 [7] I. Uber Technologies, Pyro deep universal probabilistic programming,
103 <http://pyro.ai>, accessed 2017-07-31 (2017-2018).
- 104 [8] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin,
105 S. Ghemawat, G. Irving, M. Isard, et al., Tensorflow: a system for large-
106 scale machine learning., in: OSDI, Vol. 16, 2016, pp. 265–283.
- 107 [9] S. v. d. Walt, S. C. Colbert, G. Varoquaux, The numpy array: a structure
108 for efficient numerical computation, Computing in Science & Engineering
109 13 (2) (2011) 22–30.

110 **Required Metadata**

111 **Current executable software version**

112 Ancillary data table required for sub version of the executable software:
 113 (x.1, x.2 etc.) kindly replace examples in right column with the correct
 114 information about your executables, and leave the left column as it is.

Nr.	(executable) Software metadata description	Please fill in this column
S1	Current software version	for example 1.1, 2.4 etc.
S2	Permanent link to executables of this version	example: <code>https://github.com/combogenomics/DuctApe/releases/tag/DuctApe-0.16.4</code> :
S3	Legal Software License	List one of the approved licenses
S4	Computing platform/Operating System	for example Android, BSD, iOS, Linux, OS X, Microsoft Windows, Unix-like , IBM z/OS, distributed/web based etc.
S5	Installation requirements & dependencies	
S6	If available, link to user manual - if formally published include a reference to the publication in the reference list	Example: <code>http://mozart.github.io/documentation/</code> :
S7	Support email for questions	

Table 1: Software metadata (optional)

115 **Current code version**

116 Ancillary data table required for subversion of the codebase. Kindly re-
 117 place examples in right column with the correct information about your cur-
 118 rent code, and leave the left column as it is.

Nr.	Code metadata description	Please fill in this column
C1	Current code version	For example v42
C2	Permanent link to code/repository used of this code version	For example: <i>https</i> : <i>//github.com/mozart/mozart2</i>
C3	Legal Code License	List one of the approved licenses
C4	Code versioning system used	For example svn, git, mercurial, etc. put none if none
C5	Software code languages, tools, and services used	For example c++, python, r, etc.
C6	Compilation requirements, operating environments & dependencies	
C7	If available Link to developer documentation/manual	For example: <i>http</i> : <i>//mozart.github.io/documentation/</i>
C8	Support email for questions	

Table 2: Code metadata (mandatory)