Towards representation agnostic probabilistic programming

OLE FENSKE*, MAXIMILIAN POPKO*, SEBASTIAN BADER, and THOMAS KIRSTE, Institute for Visual and Analytic Computing, Germany

Current probabilistic programming languages and tools tightly couple model representations with specific inference algorithms, preventing experimentation with novel representations or mixed discrete-continuous models. We introduce a factor abstraction with five fundamental operations that serve as a universal interface for manipulating factors regardless of their underlying representation. This enables representation-agnostic probabilistic programming where users can freely mix different representations (e.g. discrete tables, Gaussians distributions, sample-based approaches) within a single unified framework, allowing practical inference in complex hybrid models that current toolkits cannot adequately express.

1 Introduction

Probabilistic programming languages (PPLs) and toolkits (PPTs) enable practitioners to express complex statistical models and perform Bayesian inference without manually implementing inference algorithms. However, the scope of models that can be defined successfully depends on the mechanisms available for representing distributions. It is interesting to note that several PPLs (e. g. PyMC [Salvatier et al. 2016], Stan [Carpenter et al. 2017], Pyro [Bingham et al. 2019]) tightly couple model representations with specific inference algorithms (e.g., the NUTS sampler for Stan). The PPTs Factorie [McCallum et al. 2009], RXInfer [Bagaev et al. 2023], Infernet [Minka et al. 2018] provide a wider range of representations, based on the concept of "factors", but do not offer a standard way to define new representations.

However, efficient computational probabilistic reasoning often demands highly specialized representations. For instance, structured probability spaces [Choi et al. 2015] – distributions over scattered sets – are difficult to represent using "standard" representations based on samples, (sparse) arrays, or parametric mechanisms; instead they rely on probabilistic sentential decision diagrams [Kisa et al. 2014]. For a "universal" probabilistic reasoning library, it would therefore be desirable to be able to add such specialized representations. In addition, one would like the library API to provide a set of operations that allows to formulate all probabilistic computations in a way that is agnostic to the specific representation chosen.

The factor concept [Koller and Friedman 2009], which generalizes the concept of distribution functions, is in principle a powerful abstraction tool. Combined with a standard set of factor operations, it is able to decouple model syntax (probabilistic structure) from semantics (computational realization). This allows inference algorithms, like filtering or smoothing, to be formulated as factor expressions

independent of the underlying representation. We therefore advocate to provide the "factor" as fundamental abstraction and remove any assumptions about concrete representations of factors in the factor-level API provided by the toolkit.

In this paper, we outline the basic structure of the factor-level API of such a tool and give an intuitive example utilizing this formalism.

2 Factors and factor operations

A factor f_{XY} over random variables X and Y is an abstract mathematical function $f : X \times Y \rightarrow \mathbb{R}$ that assigns a real number to every configuration (x, y). Factors generalize familiar concepts: probability tables for discrete variables, Gaussian distributions for continuous variables, mixed discrete-continuous (conditional Gaussians) [Lauritzen 1992] and unnormalized potentials are all factors. Crucially, the abstract mathematical definition is intentionally separate from how a factor is represented in a computer.

Five fundamental operations serve as the API for manipulating factors regardless of their representation:

Table 1. Factor operations.

Operation	Expression	Usage
Multiplication	$h_{XYZ} = f_{XY} \otimes g_{YZ}$	Joining
Sum-Out	$g_Y = \oint_X f_{XY}$	Marginalization
Reduction	$g_Y = f_{XY}^{(X=x)}$	Conditioning
Division	$h_{XYZ} = f_{XY} \oslash g_{YZ}$	Smoothing
Addition	$h_X = f_X \oplus g_X$	Mixture

The first three operations are the conventional standard factor operations that form the algebraic core of message passing in factor graphs (e.g. the sum-product algorithm). *Multiplication* combines messages, *Sum-Out* marginalizes variables, and *Reduction* incorporates evidence [Kschischang et al. 2001]. *Division* is a standard operation for probabilistic inference (e.g. smoothing densities[Koller and Friedman 2009]), but not always provided in PPTs (its result is not normalizable in general). *Addition* – usually only implicitly used inside a PPT – is necessary to compute mixtures of distributions and specifically for realizing the Sum-Out operation when constructing hierarchical factor representations.

We here concentrate on identifying the core set of operations required for computations that take distributions as input and return distributions as output; this constitutes the probability-theoretic core. There exists a wide range of useful operations that consume distributions but produce other mathematical objects – such as MAP estimates, mutual information, or entropy – which would naturally be part of a more comprehensive "standard library". The design and formalization of such derived operations is outside the scope of this paper. This separation reflects a semantic distinction:

---

*Both authors contributed equally to this research.

Authors' Contact Information: Ole Fenske, ole.fenske@uni-rostock.de; Maximilian Popko, maximilian.popko@uni-rostock.de; Sebastian Bader, sebastian.bader@uni-rostock.de; Thomas Kirste, thomas.kirste@uni-rostock.de, Institute for Visual and Analytic Computing, Rostock, Germany.

core operations are closed under distributions, whereas derived operations compute functionals or optimizers of distributions.

3 Representation agnosticism

Factors can use diverse representations: discrete factors might be arrays, hash-tables, or sentential decision diagrams; continuous factors can be parametric (e.g., Gaussian mean/variance) or non-parametric (samples). The separation between syntax and semantics ensures extensibility: new domain-specific representations can be added by implementing only the representation-specific methods for the factor operators, without altering the core framework.

Inference algorithms can then be implemented at the level of the factor operations introduced above (being defined as generic functions at the API level). Variable elimination [Koller and Friedman 2009] and belief propagation [Pearl 1988] for example naturally decompose into sequences of such factor operations and are explicitly designed for probabilistic inference in factor graphs. For example, filtering in state space models involves the factor expressions:

(1) **Prediction step**: $p_{X_{t+1}|y_{1:t}} = \oint_{x_t} p_{X_{t+1}|X_t} \otimes p_{X_t|y_{1:t}}$

(2) **Correction step** (based on new observation y_{t+1}):

$$\phi_{X_{t+1},y_{1:t+1}} = p_{X_{t+1}|y_{1:t}} \otimes p_{Y_{t+1}|X_{t+1}}^{(Y_{t+1}=y_{t+1})}$$

$$p_{y_{t+1}|y_{1:t}} = \oint_{x_{t+1}} \phi_{X_{t+1},y_{1:t+1}}$$

$$p_{X_{t+1}|y_{1:t+1}} = \phi_{X_{t+1},y_{1:t+1}} \oslash p_{y_{t+1}|y_{1:t}}$$

(Uppercase letters in subscripts define a factor's variables, lowercase letters are simply part of the factor name. Note that $p_{y_{t+1}|y_{1:t}}$ is factor with an empty variable list: a simple scalar.)

Here, each operation dispatches to representation-specific implementations. This maintains a homomorphism between the defined factor algebra and representation algebra and enables mixing representations within the same model (e.g. discrete factors as finite maps, continuous factors as Gaussians or samples, or hybrid factors as nested representations) and allows at the same time for seamless extension.

## 4 Exploiting representation agnosticism – a toy example

Fig. 1. (Left) Graphical representation of dependencies of the model and (Right) an illustration of the quadrants and their linear transition models.

As a simple example, consider a 2D world partitioned into four quadrants, each quadrant having a primary motion direction, as shown in Fig. 1. An object moving in this world will move according
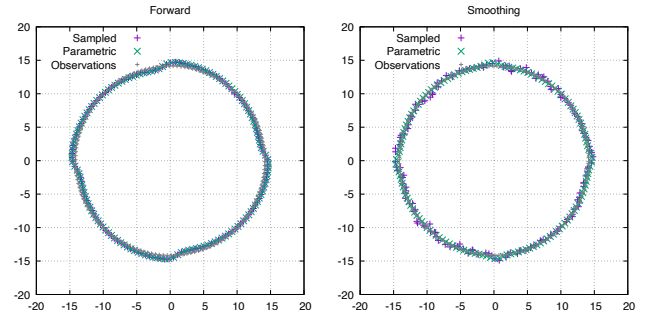

Fig. 2. Sampling versus parametric based representation. The mean marginal $F_{1:T}$ state trajectory given circular observations.

to the direction of the quadrant it is in (plus some Gaussian noise). As soon as it crosses the quadrant boundary, it will change motion direction according to the quadrant being entered. This is a simple hybrid dynamic model where a continuous variable F_t (position, range $\mathbb{R}^2$) interacts with a discrete state S_t (range $\{0, 1, 2, 3\}$, one of four quadrants in the 2D plane). Figure 1 illustrates the factorized structure of the quadrant-model, which captures the interaction between the continuous (subsymbolic) state F_t and the discrete (symbolic) state S_t over time. The continuous dynamics $p_{F_t|F_{t-1},S_{t-1}}$ thus describes smooth motion within a region, while the discrete dynamics $p_{S_t|S_{t-1},F_t}$ captures the event of crossing into the next quadrant. The observation factor $p_{Y_t|F_t}$ links the latent process to the measured data. Together, these factors define the full probabilistic structure of the model. The coupling of continuous evolution and discrete transitions yields a simple hybrid system illustrating how the factorized representation unifies symbolic reasoning with continuous-state estimation: the factor expressions that describe the probabilistic computations are invariant to the kind of random variables involved and the representation chosen.

The main computational challenge arises during the *Multiplication* ($\otimes$) of the conditional Gaussian transition factor $p_{F_t|F_{t-1},S_{t-1}}$ with the discontinuous link factor $p_{S_t|S_{t-1},F_t}$. Such a model can not be represented in a simple parametric way. A model developer might be tempted to compare different options for approximate representations – for instance, comparing sample-based representations with parametric approximations using truncated Gaussians and moments matching.

The *sampling-based representation* approximates these operations via Monte Carlo estimation or particle-based updates. The *parametric representation* expresses factors in closed form (e.g., truncated Gaussian and moments matching) and performs all operations analytically. Both use the same set of generic factor operations, but provide different methods defining their computational realization.

This separation ensures that the probabilistic model can remain fixed, while enabling flexible exploration of different factor representations and inference schemes within a single unified framework. Figure 2 shows the filter and smoothing trajectories for both parametric and sampled representations, achieved by simply swapping the underlying factor definition.

5 Conclusion

If factor expressions are defined using representation-agnostic generic operations, it becomes possible: (1) to *combine* different representations in the same model, (2) to *exchange* different representations for experimentation, and (3) to add new specialized representations – without needing to modify a given model.

Our point here is *not* that existing PPLs/PPTs lack object-oriented APIs for probabilistic computations, but rather that *representation agnosticism* at the level of model formulation seemingly has not been an explicit design objective so far – an objective that we argue should be adopted in future developments.

By treating factors and factor operations as first-class abstractions, we achieve representation-agnostic probabilistic programming where inference algorithms work uniformly across heterogeneous representations. Users can freely experiment with different representation strategies, extend the system with domain-specific representations within a unified framework. This approach enables practical inference in complex hybrid models that current toolkits cannot adequately express or efficiently solve.

As an outlook, it is interesting to consider how computations on distributions relate to sampling-based probabilistic programs. Executing a probabilistic program yields realizations drawn from a (typically joint) distribution. If such realizations are viewed as the dynamic semantics of the program, then the distribution from which they are drawn can be regarded as its static semantics, or semantic type. Under this interpretation, the operations identified in this work form a minimal core for computing the type of a probabilistic program. Recent work such as [Faggian et al. 2024] explores the construction of probabilistic program types using typed lambda calculi; our approach connects to this line of research by providing a flexible, distribution-centered foundation for computing and representing such types.

References

Dmitry Bagaev, Albert Podusenko, and Bert de Vries. 2023. RxInfer: A Julia package for reactive real-time Bayesian inference. *Journal of Open Source Software* 8, 84 (2023), 5161. doi:10.21105/joss.05161

Eli Bingham, Jonathan P. Chen, Martin Jankowiak, Fritz Obermeyer, Neeraj Pradhan, Theofanis Karaletsos, Rohit Singh, Paul Szerlip, Paul Horsfall, and Noah D. Goodman. 2019. Pyro: deep universal probabilistic programming. 20, 1 (2019), 973–978.

Bob Carpenter, Andrew Gelman, Matthew D. Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. 2017. *Stan* : A Probabilistic Programming Language. 76, 1 (2017). doi:10.18637/jss.v076.i01

Arthur Choi, Guy Van den Broeck, and Adnan Darwiche. 2015. Tractable Learning for Structured Probability Spaces: A Case Study in Learning Preference Distributions. https://www.semanticscholar.org/paper/Tractable-Learning-for-Structured-Probability-A-in-Choi-Broeck/72bb6887e437942f829822cd9090f63672dd3441

Claudia Faggian, Daniele Pautasso, and Gabriele Vanoni. 2024. Higher Order Bayesian Networks, Exactly. *Proc. ACM Program. Lang.* 8, POPL (Jan. 2024), 84:2514–84:2546. doi:10.1145/3632926

Doga Kisa, Guy Van den Broeck, Arthur Choi, and Adnan Darwiche. 2014. Probabilistic sentential decision diagrams. In *Proceedings of the Fourteenth International Conference on Principles of Knowledge Representation and Reasoning* (Vienna, Austria, 2014-07-20) *(KR'14)*. AAAI Press, 558–567.

Daphne Koller and Nir Friedman. 2009. *Probabilistic Graphical Models: Principles and Techniques - Adaptive Computation and Machine Learning.* The MIT Press.

F.R. Kschischang, B.J. Frey, and H.-A. Loeliger. 2001. Factor graphs and the sum-product algorithm. *IEEE Transactions on Information Theory* 47, 2 (2001), 498–519. doi:10.1109/18.910572

Steffen L. Lauritzen. 1992. Propagation of Probabilities, Means, and Variances in Mixed Graphical Association Models. *J. Amer. Statist. Assoc.* 87, 420 (1992), 1098–1108. arXiv:https://www.tandfonline.com/doi/pdf/10.1080/01621459.1992.10476265 doi:10.1080/01621459.1992.10476265

Andrew McCallum, Karl Schultz, and Sameer Singh. 2009. FACTORIE: Probabilistic Programming via Imperatively Defined Factor Graphs. In *Neural Information Processing Systems (NIPS)*.

T. Minka, J.M. Winn, J.P. Guiver, Y. Zaykov, D. Fabian, and J. Bronskill. 2018. /Infer.NET 0.3. Microsoft Research Cambridge. http://dotnet.github.io/infer.

Judea Pearl. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.

John Salvatier, Thomas V. Wiecki, and Christopher Fonnesbeck. 2016. Probabilistic programming in Python using PyMC3. 2 (2016), e55. doi:10.7717/peerj-cs.55 Publisher: PeerJ Inc..