

OpenGM: A C++ Library for Discrete Graphical Models

Bjoern Andres*

Thorsten Beier*

Jörg H. Kappes*

BANDRES@SEAS.HARVARD.EDU

THORSTEN.BEIER@IWR.UNI-HEIDELBERG.DE

KAPPES@MATH.UNI-HEIDELBERG.DE

HCI, University of Heidelberg, Speyerer Str. 6, 69126 Heidelberg, Germany

<http://hci.iwr.uni-heidelberg.de/opengm2>

Abstract

OpenGM is a C++ template library for defining discrete graphical models and performing inference on these models, using a wide range of state-of-the-art algorithms. No restrictions are imposed on the factor graph to allow for higher-order factors and arbitrary neighborhood structures. Large models with repetitive structure are handled efficiently because (i) functions that occur repeatedly need to be stored only once, and (ii) distinct functions can be implemented differently, using different encodings alongside each other in the same model. Several parametric functions (e.g. metrics), sparse and dense value tables are provided and so is an interface for custom C++ code. Algorithms are separated by design from the representation of graphical models and are easily exchangeable. OpenGM, its algorithms, HDF5 file format and command line tools are modular and extendible.

Keywords: Graphical Model, Combinatorial Optimization, Inference, C++

1. Introduction and Related Work

Graphical models have become a standard tool in machine learning, and inference (marginal and MAP estimation) is the central problem, cf. Nowozin and Lampert (2011).

These models can be defined rigorously as models of functions that factorize w.r.t. an associative and commutative operation, cf. Werner (2008). The C++ library OpenGM is based on this general definition that allows for a unified treatment of accumulative operations on such functions, including optimization, summation (marginalization), conjunction and disjunction. It provides a variety of inference algorithms¹ beyond message passing (Fig. 1). It can deal efficiently with large scale problems, since (i) functions that occur repeatedly need to be stored only once and (ii) when functions require different parametric or non-parametric encodings, multiple encodings can be used alongside each other, in the same model. No restrictions are imposed on the factor graph and the operations of the model, and the file format handles user extensions automatically. Furthermore, OpenGM is a template library in which elementary data types can be chosen to maximize efficiency.

Existing libraries do not have all of these properties. MRF-lib (Szeliski et al., 2008) is restricted to the min-sum semi-ring and second-order grid graphs. While it is highly efficient on these models, it is also specialized to these and not easily extendible. In contrast,

*. Authors contributed equally.

1. Not all algorithms can be used with each semi-ring.

Message passing	Graph cut	Search	Sampling	LP / ILP
Loopy BP	α -expansion	ICM	Gibbs	Dual decomposition
TRBP	$\alpha\beta$ -swap	LazyFlipper	Swendsen-Wang	Branch & cut
TRW-S	QPBO	LOC		A*

Figure 1: Algorithms provided by OpenGM: Loopy BP (Pearl, 1988; Kschischang et al., 2001), TRBP (Wainwright and Jordan, 2008), TRW-S (Kolmogorov, 2006), α -expansion, $\alpha\beta$ -swap (Boykov et al., 2001), QPBO (Rother et al., 2007), ICM (Besag, 1986), Lazy Flipper (Andres et al., 2010), LOC (Jung et al., 2009), Swendsen-Wang sampling (Barbu and Zhu, 2005), Dualdecomposition (subgradient and bundle-methods) (Kappes et al., 2012; Komodakis et al., 2011), native LP and Branch & cut using IBM ILOG Cplex, A* (Bergtholdt et al., 2010).

libDAI (Mooij, 2010) supports max-product and sum-product semi-rings which are hard-coded. The main drawback of libDAI is that it supports only dense value tables to encode functions which becomes prohibitive for models with many labels and higher-order factors. Similar to libDAI, FastInf (Jaimovich et al., 2010) focuses on message passing and does not impose any restrictions on the factor graph. In contrast to libDAI, it supports shared functions and different function types in a so-called relational model that is similar in spirit to the design of OpenGM. However, FastInf supports only sum-product semi-rings and, unlike OpenGM, has no generic template abstraction of semi-rings. The recently published library grante (Nowozin, 2012) provides shared functions and different function types. Furthermore, it comes with a set of learning methods. Unlike OpenGM, it is not template based, limited in its inference methods and published under a proprietary license.

The generality of OpenGM comes at the cost of performance. And yet, OpenGM is only slightly slower than libDAI when running loopy belief propagation on a grid graph. The highly optimized code of MRF-LIB is twice as fast for general second-order factors and 20 times as fast for standard metrics.

OpenGM is modular and extendible. The graphical model data structure, inference algorithms and different encodings of functions interoperate through well-defined interfaces.

2. Mathematical Foundation

OpenGM is built on a rigorous definition of the syntax and semantics of a graphical model. The syntax determines a class of functions that factorize w.r.t. an associative and commutative operation. In a probabilistic model, it determines the conditional independence assumptions. The semantics specify the operation and one function out of the class of all function that are consistent with the syntax.

The *syntax* (Fig. 2a) consists of a factor graph, i.e. a bipartite graph (V, F, E) , a linear order $<$ in V , a set I whose elements are called *function identifiers*, and a mapping $\gamma : F \rightarrow I$ that assigns one function identifier to each factor such that only factors that are connected to the same number of variables can be mapped to the same function identifier.

For any $v \in V$ and $f \in F$, the *factor* f is said to *depend* on the *variable* v iff $(v, f) \in E$. $\mathcal{N}(f)$ denotes the set of all variables on which f depends and $(v_j^{(f)})_{j \in \{1, \dots, |\mathcal{N}(f)|\}}$ the sequence

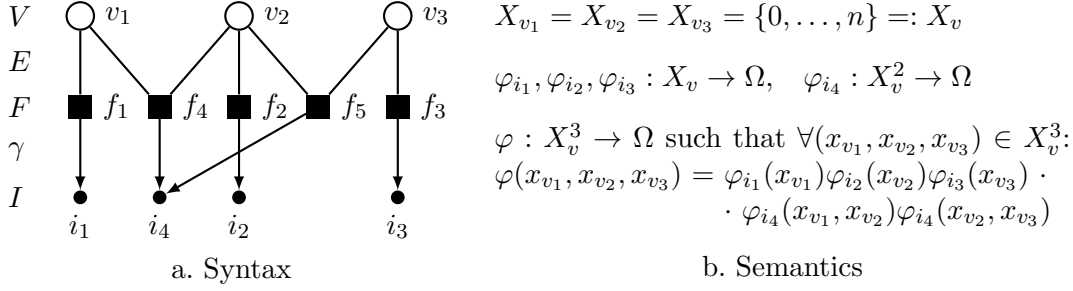


Figure 2: A factor graph (V, F, E) describes how a function φ decomposes into a product of functions. In OpenGM, we extend this syntax by a set I of *function identifiers* and a mapping $\gamma : F \rightarrow I$ that assigns one function identifier to each factor. In the above example, the factors f_4 and f_5 are mapped to the same function identifier i_4 , indicating that the corresponding functions $\varphi_{i_4}(x_{v_1}, x_{v_2})$ and $\varphi_{i_4}(x_{v_2}, x_{v_3})$ are identical.

of these variables in ascending order. Similarly, $(v_j)_{j \in \{1, \dots, |V|\}}$ denotes the sequence of *all* variables in ascending order.

Semantics (Fig. 2b) w.r.t. a given syntax consist of one finite set $X_v \neq \emptyset$ for each $v \in V$, a commutative monoid $(\Omega, \odot, 1)$ and for any $i \in I$ for which there exists an $f \in F$ with $\gamma(f) = i$, one function² $\varphi_i : X_{v_1^{(f)}} \times \dots \times X_{v_{|\mathcal{N}(f)|}^{(f)}} \rightarrow \Omega$.

The function from $X := X_{v_1} \times \dots \times X_{v_{|V|}}$ to Ω *induced* by syntax and semantics is the function $\varphi : X \rightarrow \Omega$ such that $\forall (x_{v_1}, \dots, x_{v_{|V|}}) \in X$:

$$\varphi(x_{v_1}, \dots, x_{v_{|V|}}) := \bigodot_{f \in F} \varphi_{\gamma(f)} \left(x_{v_1^{(f)}}, \dots, x_{v_{|\mathcal{N}(f)|}^{(f)}} \right). \quad (1)$$

W.l.o.g., OpenGM simplifies the syntax and semantics by substituting $V = \{0, \dots, |V| - 1\}$, equipped with the natural order, $F = \{0, \dots, |F| - 1\}$, $I = \{0, \dots, |I| - 1\}$ and for each $v \in V$, $X_v = \{0, \dots, |X_v| - 1\}$. A graphical model is thus completely defined by the number of variables $|V|$, the number of labels $|X_v|$ of each variable $v \in V$, the edges E of the factor graph, the number of functions $|I|$, the assignment of functions to factors γ , the commutative monoid $(\Omega, \odot, 1)$ and one function φ_i for each function identifier $i \in I$.

Given a graphical model and, instead of just the commutative monoid $(\Omega, \odot, 1)$, a commutative semi-ring $(\Omega, \odot, 1, \oplus, 0)$, the problem of computing

$$\bigoplus_{x \in X} \varphi(x) \quad \text{i.e.} \quad \bigoplus_{x \in X} \bigodot_{f \in F} \varphi_{\gamma(f)} \left(x_{v_1^{(f)}}, \dots, x_{v_{|\mathcal{N}(f)|}^{(f)}} \right) \quad (2)$$

is a central problem in machine learning with instances in optimization $(\mathbb{R}, +, 0, \min, \infty)$, marginalization $(\mathbb{R}^+, \cdot, 1, +, 0)$ and constrained satisfaction $(\{0, 1\}, \wedge, 1, \vee, 0)$.

3. Using and Extending OpenGM

The first step when using OpenGM is to construct a *label space* that determines the number of variables and the number of labels of each variable. The next step is to fix the data type

2. The existence of φ implies $\forall f, f' \in F : \gamma(f) = \gamma(f') \Rightarrow \forall j \in \{1, \dots, \deg(f)\} : X_{v_j^{(f)}} = X_{v_j^{(f)'}}$.

of the domain Ω , the operation \odot and the way functions are encoded, by choosing the parameters of the *graphical model* class template as in the example below.

To define a *function* such as $\varphi_{i_4}(y_1, y_2)$, one needs to indicate how many labels y_1 and y_2 have and set the parameters of the function or fill its value table. Once a function has been added to the model, it can be connected to several *factors* and thus assigned to different sets of variables. This procedure is always the same, regardless of the number and type of classes used to encode functions. Details are described in the users' section of the manual.

Algorithms for optimization and inference are classes in OpenGM. To run an algorithm, one instantiates an object of the class, providing a model and optional control parameters, and calls the member function *infer*, either without parameters or with one parameter indicating a *visitor* (see example). Visitors are a powerful tool for monitoring and controlling algorithms by code injection. Once an algorithm has terminated, results such as optima and bounds can be obtained via member functions. Detailed instructions can be found in the users' section of the manual.

OpenGM provides interfaces for custom algorithms, custom parametric functions, custom discrete spaces and custom semi-rings. These interfaces are described in the developers' section of the manual.

```

1  typedef SimpleDiscreteSpace<size_t, size_t> Space;
2  Space space(numberOfVariables, numberOfLabels);
3  typedef OPENGM_TYPELIST_2(ExplicitFunction<float>, PottsFunction<float>) Functions;
4  typedef GraphicalModel<float, Adder, Functions, Space> Model;
5  Model gm(space);
6  ExplicitFunction<float> f1(&numberOfLabels, &numberOfLabels + 1);
7  f1(0) = ...; f1(1) = ...; ...
8  Model::FunctionIdentifier fid1 = gm.addFunction(f1);
9  const size_t variableIndex = 0;
10 gm.addFactor(fid1, &variableIndex, &variableIndex + 1);
11 PottsFunction<float> f2(numberOfLabels, numberOfLabels, 0.0f, 0.3f);
12 Model::FunctionIdentifier fid2 = gm.addFunction(f2);
13 size_t variableIndices[] = {variableIndex, variableIndex + 1};
14 gm.addFactor(fid2, variableIndices, variableIndices + 2);
15 typedef BpUpdateRules<Model, Minimizer> UpdateRules;
16 typedef MessagePassing<Model, Minimizer, UpdateRules, MaxDistance> BeliefPropagation;
17 BeliefPropagation::Parameter parameter(maxNumberOfIterations, convergenceBound, damping);
18 BeliefPropagation bp(gm, parameter);
19 MessagePassingVerboseVisitor<BeliefPropagation> visitor;
20 bp.infer(visitor);
21 vector<size_t> labeling(numberOfVariables);
22 bp.arg(labeling);

```

4. Conclusion

OpenGM is a C++ library for finite graphical models that provides state-of-the-art inference algorithms. It widens the range of models representable in software by allowing for arbitrary factor graphs and semi-rings and by handling models with repetitive structure efficiently. It is fast enough for prototype development even in settings where performance is paramount. The modularity and extendibility of OpenGM, its command line tools and file format have the potential to stimulate an exchange of models and algorithms.

References

- Bjoern Andres, Jörg H. Kappes, Ullrich Köthe, and Fred A. Hamprecht. The Lazy Flipper: MAP inference in higher-order graphical models by depth-limited exhaustive search. *ArXiv e-prints*, 2010. URL <http://arxiv.org/abs/1009.4102>.
- Adrian Barbu and Song-Chun Zhu. Generalizing Swendsen-Wang to sampling arbitrary posterior probabilities. *TPMAI*, 27(8):1239–1253, 2005.
- Martin Bergtholdt, Jörg H. Kappes, Stefan Schmidt, and Christoph Schnörr. A study of parts-based object class detection using complete graphs. *IJCV*, 87(1-2):93–117, 2010.
- Julian Besag. On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society B*, 48:259–302, 1986.
- Yuri Boykov, Olga Veksler, and Ramin Zabih. Fast approximate energy minimization via graph cuts. *TPAMI*, 23(11):1222–1239, 2001.
- Ariel Jaimovich, Ofer Meshi, Ian McGraw, and Gal Elidan. FastInf: An efficient approximate inference library. *JMLR*, 11:1733–1736, 2010.
- Kyomin Jung, Pushmeet Kohli, and Devavrat Shah. Local rules for global MAP: When do they work? In *NIPS*, 2009.
- Jörg H. Kappes, Bogdan Savchynskyy, and Christoph Schnörr. A bundle approach to efficient MAP-inference by Lagrangian relaxation. In *CVPR*, 2012. in press.
- Vladimir Kolmogorov. Convergent tree-reweighted message passing for energy minimization. *TPAMI*, 28(10):1568–1583, 2006.
- Nikos Komodakis, Nikos Paragios, and Georgios Tziritas. MRF energy minimization and beyond via dual decomposition. *TPAMI*, 33(3):531–552, 2011.
- Frank R. Kschischang, Brendan J. Frey, and Hans-Andrea Loeliger. Factor graphs and the sum-product algorithm. *Transactions on Information Theory*, 47:498–519, 2001.
- Joris M. Mooij. libDAI: A free and open source C++ library for discrete approximate inference in graphical models. *JMLR*, 11:2169–2173, August 2010.
- Sebastian Nowozin. grante 1.0. <http://www.nowozin.net/sebastian/grante>, 2012.
- Sebastian Nowozin and Christoph H. Lampert. Structured learning and prediction in computer vision. *Foundations and Trends in Computer Graphics and Vision*, 6(3-4):185–365, 2011.
- Judea Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann, San Francisco, CA, USA, 1988.
- Carsten Rother, Vladimir Kolmogorov, Victor Lempitsky, and Martin Szummer. Optimizing binary MRFs via extended roof duality. In *CVPR*, 2007.
- Rick Szeliski, Ramin Zabih, Daniel Scharstein, Olga Veksler, Vladimir Kolmogorov, Aseem Agarwala, Marshall Tappen, and Carsten Rother. A comparative study of energy minimization methods for markov random fields with smoothness-based priors. *TPMAI*, 30(6):1068–1080, 2008.
- Martin J. Wainwright and Michael I. Jordan. *Graphical Models, Exponential Families, and Variational Inference*. Now Publishers Inc., Hanover, MA, USA, 2008.
- Tomáš Werner. Marginal consistency: Unifying constraint propagation on commutative semirings. In *International Workshop on Preferences and Soft Constraints*, 2008.