

# OpenGM 2.0 Manual

*Users' and Developers' Section*

*OpenGM 2.0 was developed jointly by*  
**Bjoern Andres<sup>1</sup>, Thorsten Beier<sup>2</sup> and Jörg H. Kappes<sup>3</sup>**  
*at the*  
Heidelberg Collaboratory for Image Processing  
University of Heidelberg (Germany)

This manual is written for researchers and technical practitioners who are familiar with the very basics of discrete graphical models and have some experience in generic programming in C++, e.g. from using the Standard Template Library (STL).

OpenGM [hci.iwr.uni-heidelberg.de/opengm2](http://hci.iwr.uni-heidelberg.de/opengm2)  
CMake <http://www.cmake.org/>  
HDF5 <http://www.hdfgroup.org/>  
Boost <http://www.boost.org/>  
Doxygen <http://www.stack.nl/~dimitri/doxygen/>

<sup>1</sup>[bandres@seas.harvard.edu](mailto:bandres@seas.harvard.edu)

<sup>2</sup>[thorsten.beier@iwr.uni-heidelberg.de](mailto:thorsten.beier@iwr.uni-heidelberg.de)

<sup>3</sup>[kappes@math.uni-heidelberg.de](mailto:kappes@math.uni-heidelberg.de)

# Contents

<b>1</b>	<b>Getting Started</b>	<b>5</b>
1.1	What is OpenGM? . . . . .	5
1.2	Compiling and Installing OpenGM . . . . .	5
1.2.1	Linux . . . . .	6
1.2.2	Windows . . . . .	6
1.3	Installing Optional Third Party Libraries . . . . .	6
<b>2</b>	<b>Graphical Models</b>	<b>9</b>
2.1	Mathematical Foundation . . . . .	9
2.2	The Graphical Model Class Template . . . . .	9
2.3	Label Spaces . . . . .	11
2.3.1	Simple Discrete Space . . . . .	11
2.3.2	Discrete Space . . . . .	11
2.3.3	Further Reading . . . . .	12
2.4	Functions . . . . .	12
2.4.1	Explicit Function . . . . .	13
2.4.2	Sparse Function . . . . .	13
2.4.3	Potts Function . . . . .	14
2.4.4	Potts-N Function . . . . .	14
2.4.5	Generalized Potts Functions . . . . .	15
2.4.6	Absolute Label Difference . . . . .	15
2.4.7	Squared Label Difference . . . . .	15
2.4.8	Truncated Absolute Label Difference . . . . .	16
2.4.9	Truncated Squared Label Difference . . . . .	16
2.4.10	View Functions . . . . .	16
2.4.11	Further Reading . . . . .	16
2.5	Operations . . . . .	16
2.5.1	Further Reading . . . . .	16
2.6	Factors . . . . .	17
2.6.1	Independent Factors . . . . .	18
2.6.2	Factor Arithmetic . . . . .	18
2.7	Example: A Potts Model . . . . .	19
2.8	Graphical Model Interface . . . . .	20
2.9	Loading and Saving Graphical Models . . . . .	21
<b>3</b>	<b>Inference Algorithms</b>	<b>23</b>
3.1	Mathematical Foundation . . . . .	23
3.2	Inference Algorithm Interface . . . . .	23
3.3	Message Passing . . . . .	25
3.3.1	Belief Propagation . . . . .	25
3.3.2	TRBP . . . . .	26
3.4	TRW-S* . . . . .	26
3.5	Dynamic Programming . . . . .	27

3.6	Dual Decomposition	27
3.6.1	Using Sub-Gradient Methods	28
3.6.2	Using Bundle Methods*	29
3.7	A* Search	30
3.8	(Integer) Linear Programming*	30
3.9	Graph Cuts*	31
3.10	QPBO*	32
3.11	Move Making	32
3.11.1	$\alpha$ -Expansion	33
3.11.2	$\alpha\beta$ -Swap	33
3.11.3	Iterative Conditional Modes (ICM)	34
3.11.4	Lazy Flipper	34
3.11.5	LOC	35
3.12	Sampling	35
3.12.1	Gibbs	35
3.12.2	Swendsen-Wang	38
3.13	Brute Force Search	40
3.14	Wrappers around Other Libraries	40
3.14.1	libDAI*	40
3.14.1.1	Belief Propagation	41
3.14.1.2	TRBP	41
3.14.1.3	Double Loop Generalized BP	42
3.14.1.4	Fractional BP	42
3.14.1.5	TreeExpectationPropagation	42
3.14.1.6	Exact	43
3.14.1.7	Junction Tree	43
3.14.1.8	Gibbs	43
3.14.1.9	Mean Field	44
3.14.2	MRF-LIB*	44
3.14.2.1	ICM	45
3.14.2.2	$\alpha$ -Expansion	45
3.14.2.3	$\alpha\beta$ -Swap	45
3.14.2.4	LBP	46
3.14.2.5	BP-S	46
3.14.2.6	TRW-S	46
<b>4</b>	<b>Command Line Tools</b>	<b>49</b>
4.1	Performing Inference with Command Line Tools	49
4.1.1	Basic Usage	50
4.1.2	Keeping a Protocol	51
4.1.3	Example for Usage of the Command Line Tool	51
4.1.4	Reading Protocol Files in MatLab	54
4.2	Changing Model Type	54
4.3	Adding Inference Methods	55
4.4	Adding Parameters	57
4.5	Adding Return Types	60
<b>5</b>	<b>OpenGM and MATLAB</b>	<b>61</b>
<b>6</b>	<b>Extending OpenGM</b>	<b>63</b>
6.1	Label Spaces	63
6.2	Functions	63
6.3	Operations	65
6.4	Algorithms	65

<b>7</b>	<b>Examples</b>	<b>67</b>
7.1	Markov Chain . . . . .	67
7.2	Bipartite Matching . . . . .	68
7.3	Interpixel-Boundary Segmentation . . . . .	70
<b>8</b>	<b>License</b>	<b>75</b>

# Chapter 1

## Getting Started

### 1.1 What is OpenGM?

OpenGM is a C++ template library for discrete factor graph models and distributive operations on these models. It includes state-of-the-art optimization and inference algorithms<sup>1</sup> beyond message passing. OpenGM handles large models efficiently, 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, using included and custom C++ code. No restrictions are imposed on the factor graph or the operations of the model. OpenGM is modular and extendible. Elementary data types can be chosen to maximize efficiency. The graphical model data structure, inference algorithms and different encodings of functions interoperate through well-defined interfaces. The binary OpenGM file format is based on the HDF5 standard and incorporates user extensions automatically.

### 1.2 Compiling and Installing OpenGM

OpenGM is a header-only C++ template library and can be used as such by simply decompressing the archive in a desired location. However, OpenGM also comes with command line tools, tutorials and unit tests that can be built using Cross Platform Make (CMake). We recommend to use the Cmake GUI and proceed as follows:

- Select the source directory.
- Select a build directory that is different from the source directory.
- Activate the check boxes labeled “Grouped” and “Advanced” for a comprehensive and structured overview of options.
- Select from the group BUILD optional components of OpenGM, e.g. BUILD\_TESTING, BUILD\_COMMANDLINE.
- Select from the group WITH optional external libraries, e.g. WITH\_BOOST, WITH\_CPLEX. Any such libraries need to be installed as described in Section 1.3.
- Click the button labeled “Configure” and follow the instructions.
- Click the button labeled “Generate” and follow the instructions.

---

<sup>1</sup>Not all can be used for each semiring.

### 1.2.1 Linux

On a Linux operating system, CMake writes make files such that OpenGM can now be built from the command line, in the build directory, by simply entering

```
> make 1
```

Unit tests are built if the CMake option `BUILD_TESTING` is set. To run these tests and verify that OpenGM is working properly, enter

```
> make test 1
```

To re-build the Doxygen reference documentation from the source code, enter

```
> make doc 1
```

To install OpenGM in a directory that can be specified using CMake, enter

```
> make install 1
```

### 1.2.2 Windows

On a Windows system with Microsoft Visual C++ installed, CMake can generate a Microsoft Visual Studio solution file that contains, among others, a target called `ALL_BUILD`. To build all components of OpenGM selected through CMake, open the solution file in Visual Studio, right click on the target `ALL_BUILD` and hit “Build”.

## 1.3 Installing Optional Third Party Libraries

OpenGM is a header-only C++ template library that depends exclusively on the C++ standard template library and can therefore be used without installing any third party code.

There are, however, optional components of OpenGM that are built on 3rd party code. In particular, OpenGM provides wrappers around third party inference code which make it possible to use this code with models that have been built with OpenGM. Third party code and binaries are published under different, possibly more restrictive licenses than OpenGM, cf. Section 8.

For HDF5, Boost, IBM ILOG CPLEX, ConicBundle and LibDAI, it is sufficient to download a copy and add the required path to the CMake environment. MaxFlow, IBFS, QPBO, TRW-S and MRF-LIB need to be patched in order to resolve naming conflicts. For Linux, we provide a shell script that downloads the respective 3rd party code to the directory external within the OpenGM source tree and applies the necessary patches. Code that has been installed this way will be detected automatically by CMake. This script can be invoked from the command line:

```
> make externalLibs 1
```

Windows users can either use Cygwin<sup>2</sup> to run the Linux script or download the libraries and apply the patches manually. Any feature of OpenGM that depends on 3rd party code needs to be activated by setting the corresponding CMake environment variable in the group `WITH` as described in Section 1.2. Here is a complete list of 3rd party libraries that enable additional features of OpenGM. Those marked with a \* need to be patched:

<sup>2</sup><http://www.cygwin.com>

- **HDF5** [12] HDF5 is a data model, library, and file format for storing and managing data which is used in OpenGM so save graphical models and result files of the command line tool.
- **Boost** [16] The Boost graph library is used for calculation of MaxFlow in graphcut-based methods.
- **CPLEX** [15] Cplex is a commercial library for linear, quadratic and integer programming. For academics a free version exists. In OpenGM Cplex is used for LP relaxations and ILP formulations.
- **ConicBundle** [14] The Conic Bundle Library is published under the GPL2 and used for dual decomposition bundle-method.
- **MaxFlow\*** [21] The max flow implementation of Vladimir Komogorov can be used in graphcut-based methods.
- **IBFS\*** [13] The max flow implementation by Incremental Breadth-First Search of Sagi Hed can be used in graphcut-based methods.
- **QPBO\*** [22] OpenGM provides a wrapper to the original QPBO-code of Vladimir Komogorov.
- **TRWS\*** [23] OpenGM provides a wrapper to the original TRWS-code of Vladimir Komogorov with the option to copy the data into their own data structure or using a view.
- **MRF-LIB\*** [33] OpenGM provides a wrapper to the original MRF-Library with the option to copy the data into their own data structure or using a view.
- **libDAI** [28] OpenGM provides a wrapper to lib-DAI, but copy the data into their own data structure.





# Chapter 2

## Graphical Models

### 2.1 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. 2.1a) 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 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. 2.1b) 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 function<sup>1</sup>  $\varphi_i : X_{v_1^{(f)}} \times \dots \times X_{v_{|\mathcal{N}(f)|}^{(f)}} \rightarrow \Omega$ .

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

$$\varphi(x_{v_1}, \dots, x_{v_{|V|}}) := \bigodot_{f \in F} \varphi_{\gamma(f)} \left( x_{v_1^{(f)}}, \dots, x_{v_{|\mathcal{N}(f)|}^{(f)}} \right). \quad (2.1)$$

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  $\gamma$ , the commutative monoid  $(\Omega, \odot, 1)$  and one function  $\varphi_i$  for each function identifier  $i \in I$ .

### 2.2 The Graphical Model Class Template

The four parameters of the class template `GraphicalModel` define 1. the co-domain of the function, i.e. the type of values the function can attain, 2. the operation w.r.t. which the objective function factorizes (Section 2.5), 3. data types to encode functions (Section 2.4), and 4. a label space that

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

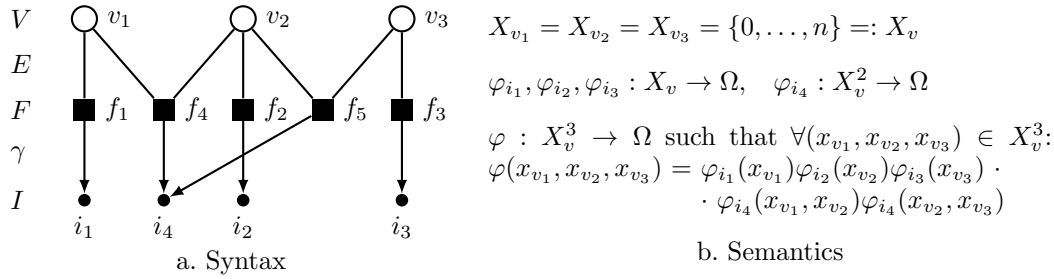


Figure 2.1: A factor graph  $(V, F, E)$  describes how a function  $\varphi$  decomposes into a product of functions. In OpenGM, we extend this syntax by a set  $I$  of *function identifiers* and a mapping  $\gamma : F \rightarrow I$  that assigns one function identifier to each factor. In the above 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.

defines a. the domain of the function, and b. the data types to index elements (variables and labels) in this domain (Section 2.3). A graphical model type is always defined like this:

```
typedef opengm::GraphicalModel<ValueType, OperationType, FunctionTypeList, SpaceType> Model; 1
```

Suitable classes for each parameter are discussed in detail in the following sections. Here are three concrete examples: First, a sum of double-valued functions which are stored explicitly, i.e. as value tables:

```
#include <opengm/graphicalmodel/graphicalmodel.hxx> 1
#include <opengm/graphicalmodel/space/discretespace.hxx> 2
#include <opengm/functions/explicit_function.hxx> 3
#include <opengm/operations/adder.hxx> 4
typedef opengm::GraphicalModel< 5
    double, 6
    opengm::Adder, 7
    opengm::ExplicitFunction<double, size_t, size_t>, 8
    opengm::DiscreteSpace<size_t, size_t> 9
> Model; 10
```

Second, a product of functions:

```
#include <opengm/graphicalmodel/graphicalmodel.hxx> 1
#include <opengm/graphicalmodel/space/discretespace.hxx> 2
#include <opengm/functions/explicit_function.hxx> 3
#include <opengm/operations/multiplier.hxx> 4
typedef opengm::GraphicalModel< 5
    double, 6
    opengm::Multiplier, 7
    opengm::ExplicitFunction<double>, 8
    opengm::DiscreteSpace<> 9
> Model; 10
```

And third, a model in which two types of functions, explicit functions and sparse functions, can be used alongside each other<sup>2</sup>:

<sup>2</sup>The class template `opengm::meta::TypeListGenerator` can assemble lists of up to 16 types. Type lists with more entries can be constructed using `opengm::meta::TypeList` in `opengm/utilities/metaprogramming.hxx`

```

#include <opengm/graphicalmodel/graphicalmodel.hxx> 1
#include <opengm/graphicalmodel/space/discretespace.hxx> 2
#include <opengm/functions/explicit.function.hxx> 3
#include <opengm/functions/sparsemarray.hxx> 4
#include <opengm/operations/multiplier.hxx> 5
#include <opengm/utilities/metaprogramming.hxx> 6
typedef opengm::meta::TypeListGenerator< 7
    opengm::ExplicitFunction<double>, 8
    opengm::SparseMarray<double> 9
>::type FunctionTypeList; 10
typedef opengm::GraphicalModel< 11
    double, 12
    opengm::Multiplier, 13
    FunctionTypeList, 14
    opengm::DiscreteSpace<> 15
> Model; 16

```

## 2.3 Label Spaces

The first step towards constructing (instantiating) a graphical model in OpenGM is to define a number of variables and, for each variable, the number of labels this variable can attain<sup>3</sup>. This is done by constructing an object called a label space.

In a space with  $n$  variables, these variables are indexed by integers ranging from 0 to  $n - 1$  as is the standard in C++. For each variable  $j$ , the labels this variable can attain are indexed in the same way, by integers from 0 to  $m_j - 1$  where  $m_j$  is the number of labels variable  $j$  can attain. All label spaces implement at least the following interface<sup>4</sup>.

```

template<class SPACE, class I = std::size_t, class L = std::size_t> 1
class SpaceBase { 2
public: 3
    typedef I IndexType; 4
    typedef L LabelType; 5
    IndexType numberOfVariables() const; 6
    IndexType numberOfLabels(const IndexType) const; 7
    bool isSimpleSpace() const; 8
    IndexType addVariable(const IndexType); 9
}; 10

```

### 2.3.1 Simple Discrete Space

In the simplest case, all variables have the same number of labels. The class template that handles this case efficiently is `SimpleDiscreteSpace`. A space with 3 variables each of which can take 4 labels is constructed like this:

```

#include <opengm/graphicalmodel/space/simplediscretespace.hxx> 1
opengm::SimpleDiscreteSpace<> space(3, 4); 2

```

### 2.3.2 Discrete Space

The most general case in which each variable can attain a different number of labels is covered by the class template `DiscreteSpace`. There are two convenient ways to construct such a space. The

<sup>3</sup>Variables that attain infinitely many labels and graphical models with infinitely many variables are currently not representable in OpenGM.

<sup>4</sup>Space class templates derive from `SpaceBase` via the curiously recurring template pattern.

first is via indices:

```
#include <opengm/graphicalmodel/space/discretespace.hxx>      1
opengm::DiscreteSpace<> space;                                  2
space.addVariable(4);                                          3
space.addVariable(2);                                          4
space.addVariable(3);                                          5
```

The second is via a pair of iterators that point to the beginning and end of a sequence of label counts:

```
#include <opengm/graphicalmodel/space/discretespace.hxx>      1
size_t numbersOfLabels[] = {4, 2, 3};                          2
opengm::DiscreteSpace<> space(numbersOfLabels, numbersOfLabels + 3); 3
```

Both examples construct a discrete space with three variables attaining 4, 2 and 3 different labels, respectively.

### 2.3.3 Further Reading

More on label spaces can be found in Section 6.1, in the reference documentation and in the header files located in `include/opengm/graphicalmodel/space`.

## 2.4 Functions

	Closed form $f(x_C) =$	Number of parameters	I/O Support
Explicit Function	$\theta_{C, x_C}$	$\prod_{c \in C}  X_c $	yes
Sparse Function	$\theta(\xi(x_C))$	$\leq \prod_{c \in C}  X_c $	yes
Potts Function	$\alpha \cdot (x_a = x_b) + \beta$	2	yes
Generalized Potts Functions	$\alpha(\gamma(x_c))$	Bell number of $ C $	yes
Multidimensional Potts Function	$\alpha \cdot \prod_{a, b \in C} (x_a = x_b) + \beta$	2	yes
Absolute Label Difference	$\alpha  x_a - x_b $	1	yes
Squared Label Difference	$\alpha (x_a - x_b)^2$	1	yes
Truncated Absolute Label Difference	$\alpha \max( x_a - x_b , T)$	2	yes
Truncated Squared Label Difference	$\alpha \max((x_a - x_b)^2, T)$	2	yes
View Function	$\alpha \cdot gm \rightarrow g(x_C)$	2	no
View Function with Offset	$\alpha \cdot gm \rightarrow g(x_C) + \theta_{x_C}$	$2 + \prod_{c \in C}  X_c $	no
Memory View Function	$\rightarrow g(x_C)$	1	no

Table 2.1: Types (classes) of functions included in OpenGM

This section describes the types (classes) of functions included in OpenGM (Tab. 2.1) and explains how functions are added to a graphical model. Once a function has been added, it can be used as often as required, with different subsets of variables, as described in Section 2.6.

All functions have at least three template parameters, `ValueType`, `IndexType` and `LabelType`. These need to be consistent with the `ValueType` of the graphical model class template and with the `IndexType` and `LabelType` of the label space, respectively. Both `IndexType` and `LabelType` are `size_t` by default therefore need not be specified.

The general interface to evaluate functions is via `operator()`. This operator expects one argument, an iterator to the beginning of a sequence of labels, one for each input variable of the function. The main advantage of this function interface is that it easily extends to higher order functions:

```

Function<double> f(...); // constructor depends on function type      1
std::vector<size_t> labels(2, 0);                                     2
double v = f(labels.begin()); // set v = f(0, 0)                     3
Function<double> g(...);                                           4
size_t labels2[] = {1, 2, 6, 3};                                    5
double w = f(labels2); // set w = f(1, 2, 6, 3)                     6

```

### 2.4.1 Explicit Function

To describe a function in terms of its full value table, OpenGM provides the class template `ExplicitFunction` that is based on an efficient implementation of runtime-flexible multi-dimensional arrays [2]. As an example, consider a function on two (so far unspecified) variables each of which assumes four different labels:

```

#include <opengm/graphicalmodel/graphicalmodel.hxx>                 1
#include <opengm/operations/adder.hxx>                               2
#include <opengm/graphicalmodel/space/simplediscretespace.hxx>      3
#include <opengm/functions/explicit_function.hxx>                    4
typedef opengm::SimpleDiscreteSpace<> Space;                         5
typedef opengm::GraphicalModel<double, opengm::Adder, opengm::ExplicitFunction<double>, Space> Model; 6
int main() {                                                         7
    Space space(10, 4);                                              8
    Model gm(space);                                                9
    size_t shape[] = {4, 4};                                        10
    opengm::ExplicitFunction<double> f(shape, shape + 2, 1.0);      11
    f(0, 0) = 0.2;                                                 12
    f(1, 1) = 0.2;                                                 13
    f(2, 2) = 0.2;                                                 14
    f(3, 3) = 0.2;                                                 15
    Model::FunctionIdentifier fid = gm.addFunction(f);              16
    return 0;                                                       17
}                                                                      18

```

Line 11 constructs the function and initializes all entries in the value table with 1.0. Lines 12–15 change four of the entries. Line 16 adds the function to the graphical model where it can hence be referred to by the function identifier `fid`.

### 2.4.2 Sparse Function

A full value table is not a parsimonious description of the function from the previous section: Among the  $4 \cdot 4 = 16$  entries in the table, 12 are identical (1.0). Therefore, it makes sense to take 1.0 to be the default value and specify only those entries that differ. The class template `SparseMarray` is designed for this purpose. The construction of a sparse function is similar to that of explicit functions, only the storage is more efficient:

```

#include <map>                                                         1
#include <opengm/graphicalmodel/graphicalmodel.hxx>                 2
#include <opengm/operations/adder.hxx>                               3
#include <opengm/graphicalmodel/space/simplediscretespace.hxx>      4
#include <opengm/functions/sparsemarray.hxx>                         5
int main() {                                                         6
    typedef opengm::SimpleDiscreteSpace<> Space;                    7
    typedef opengm::SparseMarray<std::map<size_t, double> > SparseFunction; 8
    typedef opengm::GraphicalModel<double, opengm::Adder, SparseFunction, Space> Model; 9
    Space space(10, 4);                                             10

```

```

Model gm(space); 11
size_t shape[] = {4, 4}; 12
SparseFunction f(shape, shape + 2, 1.0); 13
f(0, 0) = 0.2; 14
f(1, 1) = 0.2; 15
f(2, 2) = 0.2; 16
f(3, 3) = 0.2; 17
Model::FunctionIdentifier fid = gm.addFunction(f); 18
return 0; 19
} 20

```

### 2.4.3 Potts Function

Even a sparse function is not the most efficient way to encode the function from the previous section. In fact, the function assumes only two different values, 0.2, if the input variables take on the same label, and 1.0, otherwise. It is a `PottsFunction` that can, alternatively, be constructed like this:

```

#include <opengm/graphicalmodel/graphicalmodel.hxx> 1
#include <opengm/operations/adder.hxx> 2
#include <opengm/graphicalmodel/space/simplediscretespace.hxx> 3
#include <opengm/functions/potts.hxx> 4
typedef opengm::SimpleDiscreteSpace<> Space; 5
typedef opengm::GraphicalModel<double, opengm::Adder, opengm::PottsFunction<double>, Space> Model; 6
int main() { 7
    Space space(10, 4); 8
    Model gm(space); 9
    opengm::PottsFunction<double> f(4, 4, 0.2, 1.0); 10
    Model::FunctionIdentifier fid = gm.addFunction(f); 11
    return 0; 12
} 13

```

### 2.4.4 Potts-N Function

A `PottsNFunction` assumes one of two values, depending on whether or not *all* variables are assigned the same label. Unlike the classical Potts function, it is also defined for function of more than two variables. However, it should not be confused with the so-called  $P^N$  Potts function that is used e.g. in [20]. A `PottsNFunction` is constructed as follows:

```

#include <opengm/graphicalmodel/graphicalmodel.hxx> 1
#include <opengm/operations/adder.hxx> 2
#include <opengm/graphicalmodel/space/simplediscretespace.hxx> 3
#include <opengm/functions/pottsn.hxx> 4
typedef opengm::SimpleDiscreteSpace<> Space; 5
typedef opengm::GraphicalModel<double, opengm::Adder, opengm::PottsNFunction<double>, Space> Model; 6
int main() { 7
    Space space(10, 4); 8
    Model gm(space); 9
    const double allEqual = 0.2; 10
    const double notAllEqual = 1.0; 11
    size_t numbersOfLabels[] = {4, 5, 4}; 12
    opengm::PottsNFunction<double> f( 13
        numbersOfLabels, numbersOfLabels + 3, 14
        allEqual, notAllEqual 15
    ); 16
    Model::FunctionIdentifier fid = gm.addFunction(f); 17
}

```

```

    return 0;           18
}                       19

```

### 2.4.5 Generalized Potts Functions

A `PottsGFunction` is invariant under all permutations of labels such that, e.g.  $f(1, 1, 3, 4) = f(3, 3, 4, 1)$ . Its purpose is to assign different values to different partitions of the set of input variables, regardless which labels are used to describe these partitions. It generalizes both the `PottsFunction` and the `PottsNFunction` but requires more memory in general. As an example, consider a function that attains the value 10 if all labels are different, 11 if only the first and second are equal, 12 if only the first and third are equal, 13 if only the second and third are equal, and 14 if all are equal:

```

#include <opengm/graphicalmodel/graphicalmodel.hxx>           1
#include <opengm/operations/adder.hxx>                         2
#include <opengm/graphicalmodel/space/simplediscretespace.hxx> 3
#include <opengm/functions/pottsg.hxx>                         4
typedef opengm::SimpleDiscreteSpace<> Space;                  5
typedef opengm::GraphicalModel<double, opengm::Adder, opengm::PottsGFunction<double>, Space> Model; 6
int main() {                                                 7
    const double allEqual = 0.2;                             8
    const double notAllEqual = 1.0;                          9
    Space space(10, 4);                                     10
    Model gm(space);                                       11
    size_t numbersOfLabels[] = {4, 4, 4};                  12
    double values[5] = {10, 11, 12, 13, 14};               13
    opengm::PottsGFunction<double> f(numbersOfLabels, numbersOfLabels + 3, values); 14
    Model::FunctionIdentifier fid = gm.addFunction(f);      15
    return 0;                                             16
}                                                         17

```

### 2.4.6 Absolute Label Difference

`AbsoluteDifferenceFunction`, `SquaredDifferenceFunction`, `TruncatedAbsoluteDifferenceFunction`, and `TruncatedSquaredDifferenceFunction` are class templates that efficiently implement bi-variate functions that are used, for instance, in computer vision [34].

The absolute difference is a second order function that computes the weighted absolute difference between two labels, multiplied by a weight. The constructor takes the number of labels of the two variables and the weight:

```

#include <opengm/functions/absolute_difference.hxx>           1
opengm::AbsoluteDifferenceFunction<float> f(4, 4, 0.5);      2

```

### 2.4.7 Squared Label Difference

The squared difference is a second order function that computes the squared difference between two labels, multiplied by a weight. The constructor takes the number of labels for the two variables and the weight:

```

#include <opengm/functions/squared_difference.hxx>           1
opengm::SquaredDifferenceFunction<float> f(4, 6, 0.5);      2

```

### 2.4.8 Truncated Absolute Label Difference

The truncated absolute difference is a second order function that computes the truncated absolute difference between two labels and multiplies the result by a weight. If the absolute difference is greater than a threshold, the product of the threshold and the weight is returned. The constructor takes the number of labels of the two variables, the threshold and the weight:

```
#include <opengm/functions/truncated_absolute_difference.hxx> 1
const float threshold = 20; 2
const float weight = 0.5; 3
opengm::TruncatedSquaredDifferenceFunction<float> f(10, 10, threshold, weight); 4
```

### 2.4.9 Truncated Squared Label Difference

The truncated squared difference is a second order function that computes the truncated squared difference between the labels, multiplied by a weight. If the squared difference between the labels is greater than a threshold, the product of the threshold and the weight is returned. The constructor takes the number of labels of the two variables, the threshold and the weight:

```
#include <opengm/functions/truncated_squared_difference.hxx> 1
const float threshold = 20; 2
const float weight = 0.5; 3
opengm::TruncatedSquaredDifferenceFunction<float> f(10, 10, threshold, weight); 4
```

### 2.4.10 View Functions

The class template `ViewFunction` is used to treat a subset of one graphical model as a function in another graphical model, without copying data. `ViewConvertFunction` takes this idea one step further by allowing for arbitrary transformations of the values of the underlying graphical model which are computed only when the view function is evaluated, i.e. in a lazy fashion. `ModelViewFunction` provides memory efficient scaled views on factors, together with a piecewise offset. It facilitates efficient implementations of algorithms that depend on decompositions of a graphical model, in particular Dual Decomposition [19, 25].

View functions are used internally in inference algorithms and the developer is encouraged to use these. For modeling they are negligible.

### 2.4.11 Further Reading

More on functions can be found in the reference documentation as well as in the header files located in `include/opengm/functions`.

## 2.5 Operations

Commutative and associative operations are classes in OpenGM, equipped with member functions for the operation itself and its inverse, the neutral and inverse neutral element, and possibly an (inverse) hyper-operation and order associated with the operation (Table 2.2).

### 2.5.1 Further Reading

The interface is described in detail in Section 6.3 and in the reference documentation. The source code of the operation classes can be found in `include/opengm/operations`.



	Operation	Inverse Operation	Neutral Element	Inverse Neutral Element	Hyper-Operation	Inverse Hyper-Operation	Order
opengm::Adder	+	-	0	$\infty$	.	/	none
opengm::Multiplier	.	/	1	0	pow	root	none
opengm::Maximizer	max	min	$-\infty$	$\infty$	none	none	$\geq$
opengm::Minimizer	min	max	$\infty$	$-\infty$	none	none	$\leq$

Table 2.2: Common operations in OpenGM are Adder, Multiplier, Maximizer and Minimizer. Listed above are the corresponding inverse operations, hyper-operations and associated orders and neutral elements. Not all operations implement the full interface.

## 2.6 Factors

Factors connect functions with subsets of variables. Consider the following example, a to be modeled function  $f$  that factorizes into five factors which can be described in terms of four functions  $f_0, \dots, f_3$ :

$$f(x_0, x_1, x_2) = f_0(x_0) f_1(x_1) f_2(x_2) f_3(x_0, x_1) f_3(x_1, x_2) \quad (2.2)$$

Assume that these functions have already been constructed and added to a graphical model as described in Section 2.4, and can now be referenced by the function identifiers  $\text{fid0}, \dots, \text{fid3}$ . All that is missing to model (2.2) is to introduce factors:

```

{ size_t vi = 0; gm.addFactor(fid0, &vi, &vi + 1); } 1
{ size_t vi = 1; gm.addFactor(fid1, &vi, &vi + 1); } 2
{ size_t vi = 2; gm.addFactor(fid2, &vi, &vi + 1); } 3
{ size_t vi[] = {0, 1}; gm.addFactor(fid3, vi, vi + 2); } 4
{ size_t vi[] = {1, 2}; gm.addFactor(fid3, vi, vi + 2); } 5

```

When a factor is added it obtains a unique index that is returned by the function `addFactor`. This index can be used later to access the factor via the index operator of `opengm::GraphicalModel`:

```
const FactorType& operator[](const IndexType) const; 1
```

This operator returns a reference to a wrapper class template `opengm::Factor`. This wrapper class has a convenient interface for accessing the factor's properties. Most importantly, it is identical for all factors, regardless of the function class that is used to encode the factor.

```

template<class GRAPHICAL_MODEL> 1
class Factor { 2
public: 3
    typedef GRAPHICAL_MODEL GraphicalModelType; 4
    typedef typename GraphicalModelType::ValueType ValueType; 5
    typedef typename GraphicalModelType::LabelType LabelType; 6
    typedef typename GraphicalModelType::IndexType IndexType; 7
    typedef typename std::vector<IndexType>::const_iterator VariablesIteratorType; 8

```

```

IndexType numberOfVariables() const;          9
LabelType numberOfStates(const IndexType) const; 10
LabelType shape(const IndexType) const;       11
IndexType variableIndex(const IndexType) const; 12
ShapelteratorType shapeBegin() const;        13
ShapelteratorType shapeEnd() const;          14
VariablesIteratorType variableIndexBegin() const; 15
VariablesIteratorType variableIndexEnd() const; 16
template<class ITERATOR>                      17
    ValueType operator()(ITERATOR) const;    18
}                                             19

```

The wrapper class `Factor` is the main interface to analyze the semantics of a graphical model. It is used extensively by the algorithms described in Sections 3 and 6.4. In fact, it is the key design feature of OpenGM that allows different efficient implementations of functions to be used alongside each other while maintaining a clear separation between graphical models and algorithms that operate on these.

## 2.6.1 Independent Factors

A factor is strongly coupled to a graphical model. In fact, an instance of `opengm::Factor` cannot exist without a graphical model because the underlying function is stored only once, within the graphical model. However, it is sometimes useful to consider a factor (i.e. a function along with the indices of variables to which it is bound), independently. The class in OpenGM to facilitate this is called `opengm::IndependentFactor`. When an independent factor is constructed, usually by casting an `opengm::Factor` obtained from a graphical model by means of `operator[]` into an `opengm::IndependentFactor`, a copy of the function and the variable indices is made in the members of the `opengm::IndependentFactor` object.

## 2.6.2 Factor Arithmetic

Independent factors are amenable to factor arithmetic:

```

#include <vector>          1
#include <opengm/graphicalmodel/graphicalmodel.hxx> 2
#include <opengm/operations/adder.hxx> 3
#include <opengm/graphicalmodel/space/discretespace.hxx> 4
#include <opengm/functions/explicit_function.hxx> 5
#include <opengm/operations/adder.hxx> 6
#include <opengm/operations/minimizer.hxx> 7
typedef opengm::DiscreteSpace<> Space; 8
typedef opengm::ExplicitFunction<double> Function; 9
typedef opengm::GraphicalModel<double, opengm::Adder, Function, Space> Model; 10
typedef Model::FunctionIdentifier FunctionIdentifier; 11
typedef Model::IndependentFactorType IndependentFactor; 12
int main() { 13
    // build model 14
    size_t variableIndices[] = {3, 4, 5, 4}; 15
    Space space(variableIndices, variableIndices + 8); 16
    Model gm(space); 17
    // add functions 18
    FunctionIdentifier f0, f1; 19
    { size_t shape[] = {3, 4, 5}; Function f(shape, shape + 3); f0 = gm.addFunction(f); } 20
    { size_t shape[] = {4, 5, 4}; Function f(shape, shape + 3); f1 = gm.addFunction(f); } 21
    // add factors 22
    { size_t vi[] = {0, 1, 2}; gm.addFactor(f0, vi, vi + 3); } 23
    { size_t vi[] = {1, 2, 3}; gm.addFactor(f1, vi, vi + 3); } 24
    // simple factor arithmetic 25
    IndependentFactor a = gm[0] + gm[1]; 26
}

```

```

a += 2.0; 27
a += gm[0]; 28
opengm::Adder::op(gm[1], a); // in-place 29
IndependentFactor b; 30
opengm::Adder::op(gm[1], a, b); // not in-place 31
// accumulate over all variables 32
double minVal, maxVal; 33
std::vector<size_t> argMin; 34
a.accumulate<opengm::Minimizer>(minVal); 35
a.accumulate<opengm::Minimizer>(minVal, argMin); 36
// accumulate over subsets of variables (marginalize) 37
IndependentFactor marginal; 38
size_t accumulationIndices[] = {1, 2}; 39
a.accumulate<opengm::Minimizer>(accumulationIndices, accumulationIndices + 2, marginal); 40
return 0; 41
} 42

```

## 2.7 Example: A Potts Model

The following working example defines a Potts model on a two-dimensional grid.

```

#include <opengm/graphicalmodel/graphicalmodel.hxx> 1
#include <opengm/graphicalmodel/space/simplediscretespace.hxx> 2
#include <opengm/functions/potts.hxx> 3
#include <opengm/operations/adder.hxx> 4
using namespace std; // 'using' is used only in example code 5
using namespace opengm; 6
const size_t nx = 30; // width of the grid 7
const size_t ny = 30; // height of the grid 8
const size_t numberOfLabels = 5; 9
double lambda = 0.1; // coupling strength of the Potts model 10
// this function maps a node (x, y) in the grid to a unique variable index 11
inline size_t variableIndex(const size_t x, const size_t y) { 12
    return x + nx * y; 13
} 14
int main() { 15
    typedef SimpleDiscreteSpace<size_t, size_t> Space; 16
    Space space(nx * ny, numberOfLabels); 17
    typedef GraphicalModel<double, Adder, OPENGM_TYPELIST_2( 18
        ExplicitFunction<double>, PottsFunction<double>), Space> Model; 19
    Model gm(space); 20
    // for each node (x, y) on the grid, i.e. for each variableIndex(x, y), add 21
    // one first-order functions and one first-order factor 22
    for(size_t y = 0; y < ny; ++y) 23
        for(size_t x = 0; x < nx; ++x) { 24
            const size_t shape[] = {numberOfLabels}; 25
            ExplicitFunction<double> f(shape, shape + 1); 26
            for(size_t s = 0; s < numberOfLabels; ++s) { 27
                f(s) = (1.0 - lambda) * rand() / RAND_MAX; 28
            } 29
            Model::FunctionIdentifier fid = gm.addFunction(f); 30
            size_t variableIndices[] = {variableIndex(x, y)}; 31
            gm.addFactor(fid, variableIndices, variableIndices + 1); 32
        } 33
    // add one (!) 2nd order Potts function 34
    PottsFunction<double> f(numberOfLabels, numberOfLabels, 0.0, lambda); 35
    Model::FunctionIdentifier fid = gm.addFunction(f); 36
    // for each pair of nodes (x1, y1), (x2, y2) adjacent on the grid, 37
    // add one factor that connecting the corresponding variable indices 38

```

```

for(size_t y = 0; y < ny; ++y) 39
for(size_t x = 0; x < nx; ++x) { 40
    if(x + 1 < nx) { // (x, y) -- (x + 1, y) 41
        size_t variableIndices[] = {variableIndex(x, y), variableIndex(x + 1, y)}; 42
        sort(variableIndices, variableIndices + 2); 43
        gm.addFactor(fid, variableIndices, variableIndices + 2); 44
    } 45
    if(y + 1 < ny) { // (x, y) -- (x, y + 1) 46
        size_t variableIndices[] = {variableIndex(x, y), variableIndex(x, y + 1)}; 47
        sort(variableIndices, variableIndices + 2); 48
        gm.addFactor(fid, variableIndices, variableIndices + 2); 49
    } 50
} 51
} 52

```

## 2.8 Graphical Model Interface

The class template `GraphicalModel` inherits a rich set of member functions from the class template `FactorGraph`, the main interface for analyzing the structure (syntax) of the graphical model:

```

template<class S> 1
class FactorGraph { 2
public: 3
    typedef AccessorIterator<VariableAccessor, true> ConstVariableIterator; 4
    typedef AccessorIterator<FactorAccessor, true> ConstFactorIterator; 5

    size_t numberOfVariables() const; 6
    size_t numberOfVariables(const size_t) const; 7
    size_t numberOfFactors() const; 8
    size_t numberOfFactors(const size_t) const; 9
    size_t variableOfFactor(const size_t, const size_t) const; 10
    size_t factorOfVariable(const size_t, const size_t) const; 11

    ConstVariableIterator variablesOfFactorBegin(const size_t) const; 12
    ConstVariableIterator variablesOfFactorEnd(const size_t) const; 13
    ConstFactorIterator factorsOfVariableBegin(const size_t) const; 14
    ConstFactorIterator factorsOfVariableEnd(const size_t) const; 15
    bool variableFactorConnection(const size_t, const size_t) const; 16
    bool factorVariableConnection(const size_t, const size_t) const; 17
    bool variableVariableConnection(const size_t, const size_t) const; 18
    bool factorFactorConnection(const size_t, const size_t) const; 19
    bool isAcyclic() const; 20
    bool isConnected(marray::Vector<size_t>&) const; 21
    bool isChain(marray::Vector<size_t>&) const; 22
    bool isGrid(marray::Matrix<size_t>&) const; 23

    size_t maxFactorOrder() const; 24
    bool maxFactorOrder(const size_t) const; 25

    void variableAdjacencyMatrix(marray::Matrix<bool>&) const; 26
    void variableAdjacencyList(std::vector<std::set<size_t> >&) const; 27
    void variableAdjacencyList(std::vector<RandomAccessSet<size_t> >&) const; 28
    void factorAdjacencyList(std::vector<std::set<size_t> >&) const; 29
    void factorAdjacencyList(std::vector<RandomAccessSet<size_t> >&) const; 30
} 31

```

Arguably the most important member function of `GraphicalModel` is

```

template<class Iterator> ValueType evaluate(Iterator) const; 1

```

Its purpose is to evaluate the modeled function for a given labeling:

```

unsigned char labeling[] = {0, 1, 0};           1
Model::ValueType value = gm.evaluate(labeling); 2

```

Of course, the factors of a graphical model can also be evaluated independently, as described in Section 2.6.

## 2.9 Loading and Saving Graphical Models

One design goal of OpenGM was to simplify the exchange of graphical models by the introduction of a transparent file format. The result is a binary HDF5 format that is sufficiently general to accommodate the different ways of encoding discrete spaces and functions, including custom ones. It can be read by the OpenGM command line tools (Section 4), MATLAB, Python, R and many other environments that support HDF5.

Loading and saving graphical models is straight forward. Related functions are contained in a designated header file and namespace:

```

#include "opengm/graphicalmodel/graphicalmodel_hdf5.hxx"  1
opengm::hdf5::save(gm, "gm.h5", "toy-gm");                2
Model gm2;                                                3
opengm::hdf5::load(gm2, "gm.h5", "toy-gm");              4

```

The high-level functions that are used in OpenGM to create HDF5 files as well as to load and to save multi-dimensional arrays are contained in the header file `opengm/datastructures/marray/-marray_hdf5.hxx`. The user is encouraged to use these for their own purposes.



## Chapter 3

# Inference Algorithms

OpenGM is built around the idea of separating graphical models from algorithms that operate on these. Every algorithm is a separate class template whose control parameters are summarized in a nested class called `Parameter`. The algorithm interface facilitates the injection of custom code via visitors. Examples are verbose visitors that print intermediate states of algorithms to `std::cout` and timing visitors that measure runtime, monitor and save the progress and intermediate states of algorithms. The visitor interface is optional and comes at no runtime overhead if it is not used.

### 3.1 Mathematical Foundation

Given a graphical model (Section 2.1) and, instead of just the commutative monoid  $(\Omega, \odot, 1)$ , a commutative semi-ring  $(\Omega, \odot, 1, \oplus, 0)$ , the problem of computing

$$\bigoplus_{x \in X} \varphi(x) \quad \text{i.e.} \quad \bigoplus_{x \in X} \bigodot_{f \in F} \varphi_{\gamma(f)} \left( x_{v_1^{(f)}}, \dots, x_{v_{|\mathcal{N}(f)|}^{(f)}} \right) \quad (3.1)$$

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

$$\sum_{x \in X} \prod_{f \in F} \varphi_{\gamma(f)} \left( x_{v_1^{(f)}}, \dots, x_{v_{|\mathcal{N}(f)|}^{(f)}} \right) \quad (3.2)$$

$$(\arg) \min_{x \in X} \sum_{f \in F} \varphi_{\gamma(f)} \left( x_{v_1^{(f)}}, \dots, x_{v_{|\mathcal{N}(f)|}^{(f)}} \right) \quad (3.3)$$

$$\bigvee_{x \in X} \bigwedge_{f \in F} \varphi_{\gamma(f)} \left( x_{v_1^{(f)}}, \dots, x_{v_{|\mathcal{N}(f)|}^{(f)}} \right) \quad (3.4)$$

For any semi-ring  $(\Omega, \odot, 1, \oplus, 0)$ , marginals w.r.t. a subset  $U \subseteq V$  of fixed variables and corresponding labels  $l_u \in X_u$  for each  $u \in U$  are defined as the accumulation over the set  $X' = \{x \in X \mid \forall u \in U : x_u = l_u\}$ :

$$\bigoplus_{x \in X'} \bigodot_{f \in F} \varphi_{\gamma(f)} \left( x_{v_1^{(f)}}, \dots, x_{v_{|\mathcal{N}(f)|}^{(f)}} \right) . \quad (3.5)$$

### 3.2 Inference Algorithm Interface

In OpenGM, all inference algorithms are derived from the class template `opengm::Inference`. Its two template parameters specify 1. the graphical model that determines the operation w.r.t. which the objective function factorizes and 2. the accumulative operation used to do inference (Tab. 3.1).

	opengm::Adder	opengm::Multiplier	opengm::Or
opengm::Maximizer	energy maximization	probability maximization	-
opengm::Minimizer	energy minimization	(probability minimization)	-
opengm::Adder	-	marginalization	-
opengm::And	-	-	satisfiability

Table 3.1: Useful combination of OpenGM operations.

```

namespace opengm {
  enum InferenceTermination {
    UNKNOWN=0, NORMAL=1, TIMEOUT=2, CONVERGENCE=3, INFERENCE_ERROR=4
  };
  template <class GM, class ACC>
  class Inference {
  public:
    virtual std::string name() const = 0;
    virtual const GM& graphicalModel() const = 0;
    virtual InferenceTermination infer() = 0;
    virtual void setStartingPoint(typename std::vector<LabelType>::const_iterator);
    virtual InferenceTermination arg(std::vector<LabelType>& arg, const size_t argIndex = 1) const;
    virtual InferenceTermination args(std::vector<std::vector<LabelType> >& args) const;
    virtual InferenceTermination marginal(const size_t variableID, IndependentFactorType&) const;
    virtual InferenceTermination factorMarginal(const size_t factorID, IndependentFactorType&) const;
    virtual ValueType bound() const;
    virtual ValueType value() const;
  };
}

```

The first step towards running an inference algorithm is to construct (instantiate) an object of the corresponding class, providing, as input parameters, a graphical model and, optionally, an object of the nested class `Parameter`. Some algorithms can be initialized via the member function `setStartingPoint` that takes as input an iterator to the beginning of a sequence of labels, one for each variable in the model.

Algorithms are invoked by the member function `infer` that returns an `InferenceTermination` object indicating why inference has terminated, commonly `NORMAL`. `TIMEOUT` and `CONVERGENCE` are specializations of `NORMAL`.

Results are accessed through additional member functions. Provided that the accumulative operation is associated with a linear order of the domain  $\Omega$ , such as  $\leq$  and  $\geq$  in case of max and min, `arg` and `value` provide access to an optimal labeling and the corresponding value of the objective function. Some algorithms implement an additional member functions `arg` that provides access to the  $n$  best labelings and `bound` for a bound. If no bound is available, this function returns the neutral element of the accumulative operation. If the algorithm supports the computations of marginals (cf. (3.5) in Section 3.1), these can be accessed through the member functions `marginal` and `factorMarginal`. Marginals are provided conveniently, as independent factors (Section 2.6.1).

```

typename GraphicalModel<double, opengm::Adder> Model;
typename AnyInferenceAlgorithm<Model, opengm::Minimizer> Algorithm;
// build the model
Model gm(...);
// construct an instance of the algorithm
Algorithm::Parameter parameter;
parameter.numberOfIterations = 1000;
Algorithm algorithm(gm, parameter);
// run algorithm
algorithm.infer();
// access results

```



```

std::cout << algorithm.name() << " _has_found_the_labeling_";
std::vector<Model::LabelType> x;
algorithm.arg(x);
for(size_t j = 0; j < x.size(); ++j) {
    std::cout << x[j] << ' ';
}
std::cout << "\b_value:" << algorithm.value()
    << "\b_bound:" << algorithm.bound() << std::endl;

```

OpenGM provides a visitor framework to monitor the progress of inference algorithms. Visitor classes are taken as an input parameter of the function `infer`. They are used, for instance, by the command line tool to log upper and lower bounds over time.

Inference algorithms included in OpenGM are introduced in the following sections, together with some very basic background information and working examples which the reader is encouraged to “copy and paste”.

### 3.3 Message Passing

A large class of inference algorithms can be formulated in a message passing framework. The main operation in this framework is to calculate and propagate messages between subsets of variables. These propagations can often be understood as a re-parameterization of the original problem with the goal of establishing special properties in the re-weighted function that makes inference easier. Even if these properties correspond to fixed points of a message passing algorithm, it need not always converge to these.

In OpenGM, the message passing framework is summarized in the class template `MessagePassing`. It provides multiple message passing schedules: parallel, sequential and optimized for acyclic graphs. The template parameter `UpdateRules` determines how messages are computed. Within this framework, OpenGM implements variants of belief propagation and tree-re-weighted belief propagation. Note that both can be used with any semi-ring, always provide marginals, and afford an estimated optimal labeling if the accumulative operation is associated with a linear order.

#### 3.3.1 Belief Propagation

Belief propagation [31, 30, 26] is a message passing algorithm [26] that is applicable to a wide range of models. The implementation in OpenGM is based on [26], with optional message damping due to [30]. It has three control parameters, a maximum number of iterations, a convergence bound and a damping factor between zero and one, where zero means no damping. It can be used like this:

```

#include <opengm/inference/messagepassing.hxx>
#include <opengm/operations/maximizer.hxx>
// set up the optimizer (loopy belief propagation)
typedef BpUpdateRules<Model, Maximizer> UpdateRules;
typedef MessagePassing<Model, Maximizer, UpdateRules, MaxDistance> BeliefPropagation;
const size_t maxNumberOfIterations = 100;
const double convergenceBound = 1e-7;
const double damping = 0.0;
BeliefPropagation::Parameter parameter(maxNumberOfIterations, convergenceBound, damping);
BeliefPropagation bp(gm, parameter);
// optimize (approximately)
BeliefPropagation::VerboseVisitorType visitor;
bp.infer(visitor);
// obtain the (approximate) argmax
vector<size_t> labeling(numberOfVariables);
bp.arg(labeling);

```

The class `Maximizer` (lines 4–5) decides that the objective is to maximize the modeled function. Alternatives are `Minimizer` and `Integrator`. The class `MaxDistance` (line 5) is a metric for measuring the discrepancy between messages passed in consecutive iterations. If its value falls below the convergence bound, the algorithm stops. The visitor (line 14) injects code into the loop of the algorithm that prints intermediate states to `std::cout`. If no output is required, the visitor can be omitted in line 15 to achieve maximum performance, `bp.infer()`;

While this implementation of Belief Propagation is very general and can deal with arbitrary factor graphs and functions, significant speed-ups are achievable in more restrictive settings such as grid graphs. OpenGM makes the implementations of MRF-LIB and libDAI available through its interface. Yet another sequential variant of Belief Propagation is provided through the TRW-S interface.

### 3.3.2 TRBP

Tree Re-Weighted Belief Propagation [35] is a message passing algorithm similar to Belief Propagation. In contrast to Belief Propagation, it is based on a decomposition of the original problem into a set of acyclic sub-problems. If a re-weighting of these sub-problems exists such that the respective solutions are consistent, the original problem has been solved. However, such a re-weighting need not exist. The problem of re-weighting itself is formulated as a message passing schedule, similarly as in Belief Propagation. In addition, the expected occurrence of factors in the set of sub-problems is taken into account during the computation of messages. The set of sub-problems (subgraphs) can be passed as parameter to the algorithm or be selected automatically. In the latter case, spanning-trees are generated until all factors are covered at least once.

Marginals computed by TRBP are called pseudo-marginals and are approximations to the correct marginals. The theoretical background is described comprehensively in [35].

Our implementation of TRBP does not compute bounds on the objective. This would require additional computations for each subproblem which we have decided to avoid so far.

```

#include <opengm/inference/messagepassing.hxx> 1
#include <opengm/operations/maximizer.hxx> 2
// set up the optimizer (tree re-weighted belief propagation) 3
typedef TrbpUpdateRules<Model, Maximizer> UpdateRules; 4
typedef MessagePassing<Model, Maximizer, UpdateRules, MaxDistance> TRBP; 5
const size_t maxNumberOfIterations = 100; 6
const double convergenceBound = 1e-7; 7
const double damping = 0.0; 8
TRBP::Parameter parameter(maxNumberOfIterations, convergenceBound, damping); 9
TRBP trbp(gm, parameter); 10
// optimize (approximately) 11
TRBP::VerboseVisitorType visitor; 12
trbp.infer(visitor); 13
// obtain the (approximate) argmax 14
vector<size_t> labeling(numberOfVariables); 15
trbp.arg(labeling); 16

```

### 3.4 TRW-S\*

TRW-S [24] is a very efficient, sequential variant of TRBP that is applicable only to second order models. It is based on a decomposition of the model into acyclic graphs all of which respect a linear order on the set of variables. If two variables are connected in a subgraph, the path between both has to be monotone w.r.t. the order. This rather technical condition typically results in sparser sub-problems and allows for a very efficient implementation of block coordinate ascent in which previous computations can be re-used. In fact, TRW-S attempts to solve the LP-relaxation of the discrete problem over the local polytope. While the dual function, which is optimized, is

non-smooth, fixed points can also coincide with points from which only combinations of block-coordinate steps would lead to improvements. These are called points of weak-tree-agreement. Thanks to its special form, TRW-S computes bounds without additional effort and includes an advanced rounding scheme to obtain integer solutions.

In OpenGM, we interface the fast implementation of TRW-S by Kolmogorov [24] that is limited to the min-sum semi-ring. In addition to the explicit functions and truncated metrics provided by the original code, we implement a view to allows TRW-S to operate directly on OpenGM data structures, without the need to duplicate the model. To use this interface, one has to download and patch Kolmogorov’s original code [24] (Section 1.3).

The parameters of the TRW-S wrapper in OpenGM are: `numberOfIterations_`, a flag `useRandomStart_` to initialize messages randomly, a flag `doBPS_` to apply sequential belief propagation (BP-S) instead of TRW-S, the type of the function used by TRW-S `energyType_`, which can take the values VIEW, TABLES, TL1, TL2, and the tolerance for the relative duality gap `tolerance_` which causes TRW-S to stop if  $|E(x) - B(x)| / \max\{|E(x)|, 1\} \leq T$ .

```

#include <stdlib.h> 1
#include <opengm/graphicalmodel/graphicalmodel.hxx> 2
#include <opengm/inference/external/trws.hxx> 3
#include <opengm/operations/adder.hxx> 4
#include <opengm/operations/minimizer.hxx> 5
typedef opengm::GraphicalModel<float, opengm::Adder > GraphicalModelType; 6
typedef opengm::external::TRWS<GraphicalModelType> TRWS; 7
GraphicalModelType gm; 8
// ... 9
TRWS trws(gm); 10
trws.infer(); 11
std::cout << "value:_" << trws.value() << "_bound:_" << trws.bound() << std::endl; 12

```

## 3.5 Dynamic Programming

While Belief Propagation finds an optimal labeling for acyclic models, it effectively mimics dynamic programming and performing unnecessary computations. If one is interested only in an optimal labeling and the corresponding value of the objective function and not in any marginals, an explicit implementation of dynamic programming is preferable. OpenGM provides a solution that is currently limited to second order models.

```

#include <stdlib.h> 1
#include <opengm/graphicalmodel/graphicalmodel.hxx> 2
#include <opengm/inference/dynamicprogramming.hxx> 3
#include <opengm/operations/adder.hxx> 4
#include <opengm/operations/minimizer.hxx> 5
typedef opengm::GraphicalModel<double, opengm::Adder> GraphicalModelType; 6
typedef opengm::DynamicProgramming<GraphicalModelType, opengm::Minimizer> DP; 7
GraphicalModelType gm; 8
// ... 9
DP dp(gm); 10
dp.infer(); 11
std::cout << "value:_" << dp.value() << std::endl; 12

```

## 3.6 Dual Decomposition

OpenGM provides a rigorous implementation of the techniques of dual decomposition. The main idea is to duplicate variables in order to simplify the structure of the model. For each duplicate an additional constraint is required that ensures that both duplicates take the same state. This

constraints are in-cooperate via dual variables into the model, for a more technical formulation see [19].

After introducing duplicates each factor of the original graph is placed at least one time in the graph with the duplicates. If it is placed several times the energy is splitted. The new model often consists of several subgraphs which are not connected among each other and therefor can be optimized dependently. Typically, the construction is done in a way that the subproblems have a special form, e.g. they are acyclic an optimization becomes tractable. This is important, because optimization of the dual – which again can be understood as a re-weighting – requires optimization to calculate a subgradient at a given dual point.

The general parameters for all dual decomposition based methods are: The type of the decomposition `decompositionId_` which can be set manual `MANUAL` by the parameter `decomposition_`, a automatic decomposition into a single tree `TREE` or a set of spanning trees `SPANNINGTREES`. The maximal order of dual variables can be controlled by the parameter `maximalDualOrder_`. If it is set to 1, no dual variables are introduced for shared factors with an larger order than 1. The maximal number of dual steps can be set by the parameter `maximalNumberOfIterations_` and stopping criteria based on the absolute and relative gap can be set by `minimalAbsAccuracy_` and `minimalRelAccuracy_`, respectively.

For the updates on the dual function OpenGM provides two methods: (i) projected subgradient methods and (ii) bundle methods.

### 3.6.1 Using Sub-Gradient Methods

The naive subgradient method start from the dual point 0 and iteratively make steps into the direction of a subgradient  $s$ . The three mayor parameters to control the scale of the scaling  $\tau^t$  of the update are the stepsize-stride `stepsizeStride_` ( $\tau$ ), -scale `stepsizeScale_` ( $\beta$ ) and -exponent `stepsizeExponent_` ( $\alpha$ ), defining the stepsize

$$\tau^t = \tau \cdot \frac{1}{1 + (\beta \cdot t)^\alpha}.$$

By setting the parameter `stepsizeNormalizedSubgradient_` additional normalization by the length of the subgradient is performed such that

$$\tau^t = \tau \cdot \frac{1}{1 + (\beta \cdot t)^\alpha} \frac{1}{\|s\|_2}.$$

Under some mild conditions [19] both update rules are guarantee to converge. But practically they are to slow. Komodakis [25] suggest to use an adaptive stepsize, which can be used by setting the parameter `useAdaptiveStepsize_`, causing

$$\tau^t = \tau \cdot \frac{|\bar{E} - B|}{\|s\|_2}.$$

Since this does not work well for use, we prefer an projected adaptive stepsize, which also makes more sense from a theoretical point of view. This can be used by setting the parameter `useProjectedAdaptiveStepsize_` and cause

$$\tau^t = \tau \cdot \frac{|\bar{E} - B|}{\|s\|_2^2}.$$

Setting up a dual decomposition algorithm is a little bit more complicated, but allows to use any inference method for the subproblems. The template parameters are the graphical model type, the inference algorithm for the subproblems and the class including the dual blocks.

```
#include <stdlib.h> 1
#include <opengm/graphicalmodel/graphicalmodel.hxx> 2
#include <opengm/inference/messagepassing.hxx> 3
```

```

#include <opengm/inference/dualdecomposition/dualdecomposition_subgradient.hxx>      4
#include <opengm/operations/adder.hxx>                                             5
#include <opengm/operations/minimizer.hxx>                                         6
                                                                                   7
typedef opengm::Minimizer AccType;                                                8
typedef opengm::GraphicalModel<double, opengm::Adder> GraphicalModelType;        9
typedef opengm::DDDualVariableBlock<marray::Marray<ValueType> > DualBlockType;   10
typedef opengm::DualDecompositionBase<GraphicalModelType,DualBlockType>::SubGmType SubGmType; 11
typedef opengm::BeliefPropagationUpdateRules<SubGmType, AccType> UpdateRuleType;   12
typedef opengm::MessagePassing<SubGmType, AccType, UpdateRuleType, opengm::MaxDistance> InfType; 13
typedef opengm::DualDecompositionSubGradient<GraphicalModelType,InfType,DualBlockType> 14
DualDecompositionSubGradient;                                                    15
                                                                                   15
GraphicalModelType gm;                                                           16
// ...                                                                            17
DualDecompositionSubGradient ddsG(gm);                                           18
ddsG.infer();                                                                      19
std::cout << "value:_" << ddsG.value() << "_bound:_" << ddsG.bound() << std::endl; 20

```

### 3.6.2 Using Bundle Methods\*

As suggested in [19] OpenGM also provides an interface to use bundle methods for dual optimization. To use our bundle-solver, one has to install the Conic Bundle Library by Christoph Helmberg, available under the GPL2. After downloading the library one has to compile it and set the path in cmake and activate WITH\_BUNDLE.

The idea of the bundle-method is to build a linear approximation of the dual function around the current working point together with a trust region term, that keeps optimization local. See [19] and the references therein for technical details. While compared to subgradient methods, bundle-methods requires to solve a small QP for updating the dual, the number of iterations required is typically much less, such that overall bundle-methods performs faster. Furthermore there seemed to be more robust with respect to the choice of parameters.

Parameters for the bundle method are: The maximal size of bundle `maxBundlesize_`, upper and lower bound on the trust region weight `minDualWeight_` and `maxDualWeight_`. Furthermore, the flag `activeBoundFixing_` can be used to fix variables automatically to the center value if their bounds are strongly active, `noBundle_` to use a special solver that only employs a minimal bundle consisting of just one new and one aggregate gradient so that there is no real bundle available, `useHeuristicStepsize_` to uses heuristic for `stepsize/trustregion-radius` by Kiewel.

```

#include <stdlib.h>                                                                 1
#include <opengm/graphicalmodel/graphicalmodel.hxx>                               2
#include <opengm/inference/messagepassing.hxx>                                     3
#include <opengm/inference/dualdecomposition/dualdecomposition_bundle.hxx>         4
#include <opengm/operations/adder.hxx>                                             5
#include <opengm/operations/minimizer.hxx>                                         6
                                                                                   7
typedef opengm::Minimizer AccType;                                                8
typedef opengm::GraphicalModel<double, opengm::Adder> GraphicalModelType;        9
typedef opengm::DDDualVariableBlock<marray::Marray<ValueType> > DualBlockType;   10
typedef opengm::DualDecompositionBase<GraphicalModelType,DualBlockType>::SubGmType SubGmType; 11
typedef opengm::BeliefPropagationUpdateRules<SubGmType, AccType> UpdateRuleType;   12
typedef opengm::MessagePassing<SubGmType, AccType, UpdateRuleType, opengm::MaxDistance> InfType; 13
typedef opengm::DualDecompositionBundle<GraphicalModelType,InfType,DualBlockType> 14
DualDecompositionBundle;                                                         15
                                                                                   15
GraphicalModelType gm;                                                           16
// ...                                                                            17
DualDecompositionBundle ddb(gm);                                                 18
ddb.infer();                                                                      19
std::cout << "value:_" << ddb.value() << "_bound:_" << ddb.bound() << std::endl; 20

```

### 3.7 A\* Search

The A\* search algorithm suggested in [4, 18] is a branch-and-bound algorithm. For the calculation of the bound a tree-based approximation is used which guarantee to bound the optimal objective of a subset of states. For this highly efficient calculation of bounds the set of allowed branches is restricted and predefined by the parameters of the algorithm. While the method does not scale to large scale problems, it is state-of-the-art for models with a moderate size of variables (up to 100). The density of the model and the number of labels are from minor important. Especially when the pairwise and higher order terms are highly discriminative, the A\*-method performs fast. In general it guarantee to find the best or N-best states, but has no guarantee for polynomial runtime.

It is very important to select the node order and the tree used for the bound calculation appropriate. The version 2.0 includes no reasonable automatic selection – we hope to include this in a later release – especially if your models have a irregular topology one can gain a speed up of several magnitudes. For large models it would make sense to restrict the maximal heap size. If the maximal heap size is reached, the worst subsets are dropped. Since their bound is known, it may be still possible to guarantee optimality, see [4, 18] for details. All parameters listed below:

maxHeapSize_	maximum size of the heap
numberOfOpt_	number of N-best solutions that should be found
objectiveBound_	states that a worse are not considered as solution
heuristic_	method used to calculate bound/heuristic
→ DEFAULTHEURISTIC	...select fastest one automatically
→ FASTHEURISTIC	... use dynamic programming for second order models
→ STANDARDHEURISTIC	... use belief propagation
nodeOrder_	vector containing the nodes in the order they should be proceed
treeFactorIds_	factor ids spanning the tree used for bound calculation

```

#include <stdlib.h> 1
#include <opengm/graphicalmodel/graphicalmodel.hxx> 2
#include <opengm/inference/astar.hxx> 3
#include <opengm/operations/adder.hxx> 4
#include <opengm/operations/minimizer.hxx> 5
typedef opengm::Minimizer AccType; 6
typedef opengm::GraphicalModel<double, opengm::Adder> GraphicalModelType; 7
typedef opengm::AStar<GraphicalModelType, opengm::Minimizer> ASTAR; 8
GraphicalModelType gm; 9
// ... 10
ASTAR astar(gm); 11
astar.infer(); 12
std::cout << "value:_" << astar.value() << "_bound:_" << astar.bound() << std::endl; 13

```

### 3.8 (Integer) Linear Programming\*

Any energy minimization problem can be formulated as an (integer) linear program. Often it is also sufficient to consider a relaxation of the original problem by relaxing the system of inequalities. The common supposition that such problem formulations are generally computationally intractable is not true! OpenGM provides a wrapper that transform such problems in either an LP-relaxation with respect to the first order local polytope [35] or an integer linear programming formulation which is solved via branch-and-cut.

To use these methods, CPLEX [15] needs to be installed. Versions free for academic use can be obtained through the IBM academic initiative. Once IBM ILOG CPLEX and IBM CONCERT are installed, the respective paths need to be set in the OpenGM CMake environment and activated by checking WITH\_CPLEX. After a re-build, LP and ILP formulations and the CPLEX solver are available via the class `opengm::LPCplex`.

The corresponding parameter class takes many optional parameters which can be used to tune the solver. The most important for beginners are the flag to switch to integer mode `integerConstraint_`, the number of threads that should be used `numberOfThreads_` (0=autosect), the verbose flag `verbose_` that switches the CPLEX output on, and the time limit `timeLimit_` that stops the method after  $N$  seconds realtime.

```

#include <stdlib.h> 1
#include <opengm/graphicalmodel/graphicalmodel.hxx> 2
#include <opengm/operations/adder.hxx> 3
#include <opengm/operations/minimizer.hxx> 4
#include <opengm/inference/lpcplex.hxx> 5
typedef opengm::GraphicalModel<double, opengm::Adder> GraphicalModelType; 6
typedef LPCplex<GraphicalModelType, opengm::Minimizer> LPCPLEX; 7
GraphicalModelType gm; 8
// ... 9
//Set up LP 10
LPCPLEX lp(gm); 11
lp.infer(); 12
std::cout << "value:_" << lp.value() << "_bound:_" << lp.bound() << std::endl; 13
//Set up ILP 14
typename LPCPLEX::Parameter para; 15
para.integerConstraint_ = true; 16
para.timeLimit_ = 60.0; 17
para.verbose_ = true; 18
para.numberOfThreads_ = 4; 19
LPCPLEX ilp(gm,para); 20
ilp.infer(); 21
std::cout << "value:_" << ilp.value() << "_bound:_" << ilp.bound() << std::endl; 22

```

### 3.9 Graph Cuts\*

OpenGM comes with a framework for optimization by graph cuts [8], for graphical models whose operation is addition, whose variables are binary<sup>1</sup> and whose factors are submodular and of order at most 3. It can be used with the different st-min-cut/max-flow algorithms from the boost graph library, by Kolmogorov [21] and an implementation of *Maximum Flows By Incremental Breadth First Search* (IBFS) [10] by Hed [13]. As an example, consider a graphical model `gm` that has the necessary properties, and the following graph cut optimizer:

```

#include <iostream> 1
#include <opengm/graphicalmodel/graphicalmodel.hxx> 2
#include <opengm/operations/adder.hxx> 3
#include <opengm/inference/graphcut.hxx> 4
#include <opengm/inference/auxiliary/minstcutkolmogorov.hxx> 5
#include <opengm/operations/minimizer.hxx> 6
typedef opengm::GraphicalModel<double, opengm::Adder> GraphicalModelType; 7
typedef opengm::external::MinSTCutKolmogorov<size_t, double> MinCut; 8
typedef opengm::GraphCut<GraphicalModelType, opengm::Minimizer, MinStCutType> MinCut; 9
MinCut mincut(gm); 10
mincut.infer(); 11
std::cout << "value:_" << mincut.value() << std::endl; 12

```

The next example uses a push-relabel algorithm from the boost graph library that is designed for integer labeled graphs<sup>2</sup>. The fractional values of the graphical model are scaled by 1000000 and truncated.

<sup>1</sup>For submodular functions on variables with more than two labels, consider the move making algorithms described in Section 3.11.

<sup>2</sup>IBFS is also designed for integer-valued max-flow problems.

```

#include <iostream> 1
#include <opengm/graphicalmodel/graphicalmodel.hxx> 2
#include <opengm/operations/adder.hxx> 3
#include <opengm/inference/graphcut.hxx> 4
#include <opengm/inference/auxiliary/minstcutboost.hxx> 5
#include <opengm/operations/minimizer.hxx> 6
typedef opengm::GraphicalModel<double, opengm::Adder> GraphicalModelType; 7
typedef opengm::MinSTCutBoost<size_t, long, opengm::PUSH_RELABEL> MinStCutType; 8
typedef opengm::GraphCut<GraphicalModelType, opengm::Minimizer, MinStCutType> MinGraphCut; 9
MinGraphCut::Parameter para(1000000); 10
MinCut mincut(gm,para); 11
mincut.infer(); 12
std::cout << "value:_" << mincut.value() << std::endl; 13

```

### 3.10 QPBO\*

QPBO (quadratic pseudo-boolean optimization) is based on the principle of roof duality [11] and was proposed for unconstrained quadratic binary programming and network-flow problems by Boros and Hammer [6, 7]. Its first application to graphical models is due to Rother et al. [32]. OpenGM wraps the original code by Kolmogorov [22] and, in addition, provides its own QPBO solver that builds the corresponding max-flow problem, solves it with a graph cut solver and transform the resulting flow and cut back into a solution of the original problem. We suggest to use Kolmogorov's code which is faster.

QPBO can be applied on any second order binary model. If the original problem is permuted submodular, it is solved to optimality. In general, QPBO yields the solution of the LP relaxation over the local polytope which is half-integral, i.e. each attains a label 0, 1 or  $\frac{1}{2}$ . Furthermore, QPBO provides a partial optimal labeling for some variables, a property known as persistency [32]. Even if the complete optimal labeling remains unknown, parts of the problem can often be solved to optimality.

```

#include <iostream> 1
#include <opengm/graphicalmodel/graphicalmodel.hxx> 2
#include <opengm/operations/adder.hxx> 3
#include <opengm/inference/external/qpbo.hxx> 4
#include <opengm/operations/minimizer.hxx> 5
typedef opengm::GraphicalModel<double, opengm::Adder> GraphicalModelType; 6
typedef opengm::external::QPBO<GraphicalModelType> QPBO; 7
GraphicalModelType gm; 8
// ... 9
QPBO qpbo(gm); 10
qpbo.infer(); 11
std::vector<bool> optimalVariables; 12
double partialOpt = partialOptimality(optimalVariables); 13
std::cout << "value:_" << qpbo.value() << "_bound:_" << qpbo.bound() << std::endl; 14
std::cout << "partial_optimality:_" << partialOpt << std::endl; 15

```

### 3.11 Move Making

OpenGM provides several algorithms that attempt to improve an initial labeling by moving to better labelings via a constrained set of moves. This class of algorithms include  $\alpha$ -Expansion [8] (AlphaExpansion),  $\alpha\beta$ -Swap [8] (AlphaBetaSwap), ICM [5] (ICM), and the Lazy Flipper [1] (LazyFlipper).



### 3.11.1 $\alpha$ -Expansion

The  $\alpha$ -expansion algorithm [8] starts from an initial labeling and continues to iterates over all labels  $\alpha$  that variables can attain. In each iteration, a problem with binary variables is constructed from the original problem, based on the following question: Which subset of variables whose current label is not  $\alpha$  should be labeled  $\alpha$  to yield an optimal improvement w.r.t. the current labeling? In OpenGM, any algorithm can be used for the binary sub-problems. It should be chosen such that sub-problems are solved to optimality. Under mild conditions [8], each move can be reduced to a max-flow problem and solved efficiently. If no improvements are possible for any  $\alpha$ , the algorithm has converged.

```

#include <iostream> 1
#include <opengm/graphicalmodel/graphicalmodel.hxx> 2
#include <opengm/operations/adder.hxx> 3
#include <opengm/operations/minimizer.hxx> 4
#include <opengm/inference/graphcut.hxx> 5
#include <opengm/inference/alphaexpansion.hxx> 6
#include <opengm/inference/auxiliary/minstcutkolmogorov.hxx> 7
typedef opengm::GraphicalModel<double, opengm::Adder> GraphicalModelType; 8
typedef opengm::external::MinSTCutKolmogorov<size_t, double> MinStCutType; 9
typedef opengm::GraphCut<GraphicalModelType, opengm::Minimizer, MinStCutType> MinGraphCut; 10
typedef opengm::AlphaExpansion<GraphicalModelType, MinGraphCut> MinAlphaExpansion; 11
GraphicalModelType gm; 12
// ... 13
MinAlphaExpansion ae(gm); 14
ae.infer(); 15
std::cout << "value:." << ae.value() << std::endl; 16

```

More details on the parameters can be found in the reference documentation. For grid graphs, OpenGM provides an interface to the MRF-LIB (Section 3.14.2) which contains the original implementation of [8] that is more efficient for these types of problems.

### 3.11.2 $\alpha\beta$ -Swap

The  $\alpha\beta$ -swap algorithm [8] is similar to  $\alpha$ -expansion, except that it iterated over all pairs of distinct labels  $\alpha$  and  $\beta$  and, in each iteration, a binary problem is constructed based on the question which variables that are currently labeled  $\alpha$  should be labeled  $\beta$  such that the improvement over the current labeling is optimal. All other variables are fixed. If no  $\alpha\beta$ -swap can improve the current labeling, the algorithm has converged. Under mild conditions [8], each move can be reduced to a max-flow problem and solved efficiently. The implementation in OpenGM allows any algorithm to be used for the sub-problems. It should solve the sub-problems to optimality.

```

#include <iostream> 1
#include <opengm/graphicalmodel/graphicalmodel.hxx> 2
#include <opengm/operations/adder.hxx> 3
#include <opengm/operations/minimizer.hxx> 4
#include <opengm/inference/graphcut.hxx> 5
#include <opengm/inference/alphabetaswap.hxx> 6
#include <opengm/inference/auxiliary/minstcutkolmogorov.hxx> 7
typedef opengm::GraphicalModel<double, opengm::Adder> GraphicalModelType; 8
typedef opengm::external::MinSTCutKolmogorov<size_t, double> MinStCutType; 9
typedef opengm::GraphCut<GraphicalModelType, opengm::Minimizer, MinStCutType> MinGraphCut; 10
typedef opengm::AlphaBetaSwap<GraphicalModelType, MinGraphCut> MinAlphaBetaSwap; 11
GraphicalModelType gm; 12
// ... 13
MinAlphaBetaSwap abs(gm); 14
abs.infer(); 15
std::cout << "value:." << abs.value() << std::endl; 16

```

Further details on the parameters can be found in the reference documentation. For grid graphs, OpenGM provides an interface to the MRF-LIB (Section 3.14.2) that contains the original implementation of [8] and is more efficient for these types of problems.

### 3.11.3 Iterative Conditional Modes (ICM)

In ICM [5], one iterates over the variables. For each variable, the optimal labeling is chosen, conditioned to all other variables being fixed. The process is repeated until none change of a single label leads to an improvement<sup>3</sup>. This method is fast and simple. However, its fixed points are only guaranteed to be optimal within a Hamming distance of 1 which makes the algorithm prone to local optima and sensitive to the initialization.

```

#include <iostream> 1
#include <opengm/graphicalmodel/graphicalmodel.hxx> 2
#include <opengm/operations/adder.hxx> 3
#include <opengm/operations/minimizer.hxx> 4
#include <opengm/inference/icm.hxx> 5
typedef opengm::GraphicalModel<double, opengm::Adder> GraphicalModelType; 6
typedef opengm::ICM<GraphicalModelType, opengm::Minimizer> ICM; 7
typedef GraphicalModelType::LabelType LabelType; 8
GraphicalModelType gm; 9
// ... 10
std::vector<LabelType> startPoint; 11
//.. 12
ICM::Parameter para(startPoint); 13
ICM icm(gm,para); 14
icm.infer(); 15
std::cout << "value:." << icm.value() << std::endl; 16

```

### 3.11.4 Lazy Flipper

Similar to ICM, the Lazy Flipper tries to improve a initial labeling sequentially. Instead of considering only a single variable at a time, it takes into account also large connected subsets of variables up to a prescribed size. For a subgraph size of one, it reduce to ICM. When the Lazy Flipper converges, the labeling is guaranteed to be optimal within a Hamming distance equal to the subgraph size, i.e. no better labeling exists that can be reached from the current labeling by flipping the same number of variables or less. The algorithm keeps track of the history of visited subgraphs and successful flips.

```

#include <iostream> 1
#include <opengm/graphicalmodel/graphicalmodel.hxx> 2
#include <opengm/operations/adder.hxx> 3
#include <opengm/operations/minimizer.hxx> 4
#include <opengm/inference/lazyflipper.hxx> 5
typedef opengm::GraphicalModel<double, opengm::Adder> GraphicalModelType; 6
typedef opengm::LazyFlipper<GraphicalModel, opengm::Minimizer> LazyFlipper; 7
typedef GraphicalModelType::LabelType LabelType; 8
GraphicalModelType gm; 9
// ... 10
size_t maxSubgraphSize = 2; 11
std::vector<LabelType> startPoint; 12
//.. 13
LazyFlipper::Parameter para(maxSubgraphSize,startPoint); 14
LazyFlipper lf(gm,para); 15

```

<sup>3</sup>The implementation in OpenGM is more efficient.

```
lf.infer(); 16
std::cout << "value:_" << lf.value() << std::endl; 17
```

### 3.11.5 LOC

LOC [17] is a stochastic move making algorithm that starts from an arbitrary initial labeling. In each iteration, a connected subset of variables is selected at random. The problem is then solved exactly for this subset, keeping all other variables fixed. The parameters of the algorithm are `phi`, the parameter of the truncated geometric distribution from which a subgraph radius is drawn, `maxRadius`, the maximum such radius, `maxIteration`, the maximum number of iterations and `aStarThreshold`, the subgraph size beyond which Astar, instead of Bruteforce, is used to solve the sub-problem.

```
#include <opengm/inference/loc.hxx> 1
#include <opengm/operations/minimizer.hxx> 2
// set up the optimizer (bp) 3
typedef LOC<Model, Minimizer> LOC; 4
typedef Model::LabelType LabelType; 5
double phi = 0.5; 6
size_t maxRadius = 10; 7
size_t maxIterations = 100; 8
size_t aStarThreshold = 5; 9
LOC::Parameter parameter(phi,maxRadius,maxIterations,aStarThreshold); 10
LOC loc(gm, parameter); 11
// set starting point 12
std::vector<LabelType> startingPoint; 13
// assuming startingPoint has been filled 14
// with meaningful labels 15
loc.setStartingPoint(startingPoint.begin()); 16
// optimize (approximately) 17
loc.infer( ); 18
// obtain the (approximate) argmin 19
vector<LabelType> labeling(numberOfVariables); 20
bp.arg(labeling); 21
```

## 3.12 Sampling

OpenGM comes with two Markov Chain Monte Carlo algorithms, Gibbs sampling [9] and a generalization of Swendsen-Wang sampling [3]. Both methods are special cases of the Metropolis-Hastings algorithm [27] and can be used to approximate the joint probability mass function defined by a graphical model, to search for its modes as well as to approximate marginals of arbitrary subsets of variables. Internally, the algorithms generate correlated samples as states of a Markov chain. Each sample is a labeling of all variables. This Markov chain is defined by a proposal probability  $q(x'|x_t)$  of moving from the current (e.g. an initial) labeling  $x_t$  to a new labeling  $x'$ , or else sample the current labeling again.

### 3.12.1 Gibbs

In Gibbs sampling, the proposal probability is chosen such that at most one variable can change its label. This variable can either be drawn randomly from the uniform distribution over all variables or selected cyclicly. The new label is always drawn randomly from the uniform distribution over all labels the proposed variable can attain.

In each iteration, the proposed labeling  $x'$  is accepted depending on the ratio  $r = p(x')/p(x_t)$  where  $p(x)$  is the full probability mass function defined by the graphical model. If  $r \geq 1$ , the new

labeling is accepted, i.e.  $x_{t+1} = x'$ . Otherwise,  $x'$  is accepted with probability  $r$ . If  $x'$  is rejected, then  $x_{t+1} = x_t$ .

The class `Gibbs` can be used with graphical models in which the operation is multiplication or addition. In the case of multiplication, the class samples from the probability mass function that is proportional to the modeled function. This means that one does not need to normalize when setting up the graphical model. If the operation of the graphical model is addition and the modeled function is  $f(x)$ , the probability mass function is taken to be proportional to  $\exp(-f(x))$ .

The following example demonstrates how to search for the mode of a distribution and how to sample marginals of specific subsets of variables using the `MarginalVisitor`.

```

#include <vector> 1
#include <iostream> 2
3
#include <opengm/graphicalmodel/space/simplediscretespace.hxx> 4
#include <opengm/functions/explicit.function.hxx> 5
#include <opengm/functions/potts.hxx> 6
#include <opengm/operations/multiplier.hxx> 7
#include <opengm/operations/maximizer.hxx> 8
#include <opengm/graphicalmodel/graphicalmodel.hxx> 9
#include <opengm/inference/gibbs.hxx> 10
11
typedef opengm::SimpleDiscreteSpace<> Space; 12
typedef opengm::ExplicitFunction<double> ExplicitFunction; 13
typedef opengm::PottsFunction<double> PottsFunction; 14
typedef opengm::meta::TypeListGenerator<ExplicitFunction, PottsFunction>::type FunctionTypes; 15
typedef opengm::GraphicalModel<double, opengm::Multiplier, FunctionTypes, Space> GraphicalModel; 16
typedef opengm::Gibbs<GraphicalModel, opengm::Maximizer> Gibbs; 17
typedef opengm::GibbsMarginalVisitor<Gibbs> MarginalVisitor; 18
19
// build a model with 10 binary variables in which 20
// - the first variable is more likely to be labeled 1 than 0 21
// - neighboring variables are more likely to have similar labels than dissimilar 22
void buildGraphicalModel(GraphicalModel& gm) { 23
    const size_t numberOfVariables = 10; 24
    const size_t numberOfLabels = 2; 25
    Space space(numberOfVariables, numberOfLabels); 26
    gm = GraphicalModel(space); 27
    28
    // add 1st order function 29
    GraphicalModel::FunctionIdentifier fid1; 30
    { 31
        ExplicitFunction f(&numberOfLabels, &numberOfLabels + 1); 32
        f(0) = 0.2; 33
        f(1) = 0.8; 34
        fid1 = gm.addFunction(f); 35
    } 36
    37
    // add 2nd order function 38
    GraphicalModel::FunctionIdentifier fid2; 39
    { 40
        const double probEqual = 0.7; 41
        const double probUnequal = 0.3; 42
        PottsFunction f(2, 2, probEqual, probUnequal); 43
        fid2 = gm.addFunction(f); 44
    } 45
    46
    // add 1st order factor (at first variable) 47
    { 48
        size_t variableIndices[] = {0}; 49
        gm.addFactor(fid1, variableIndices, variableIndices + 1); 50
    } 51
    52
    // add 2nd order factors 53
    for(size_t j = 0; j < numberOfVariables - 1; ++j) { 54
        size_t variableIndices[] = {j, j+1}; 55

```

```

    gm.addFactor(fid2, variableIndices, variableIndices + 2);
  }
}
// use Gibbs sampling for the purpose of finding the most probable labeling
void gibbsSamplingForOptimization(const GraphicalModel& gm) {
  const size_t numberOfSamplingSteps = 1e4;
  const size_t numberOfBurnInSteps = 1e4;
  Gibbs::Parameter parameter(numberOfSamplingSteps, numberOfBurnInSteps);
  Gibbs gibbs(gm, parameter);
  gibbs.infer();
  std::vector<size_t> argmax;
  gibbs.arg(argmax);

  std::cout << "most_probable_labeling_sampled:\n";
  for(size_t j = 0; j < argmax.size(); ++j) {
    std::cout << argmax[j] << ",\n";
  }
  std::cout << "\b\b)" << std::endl;
}
// use Gibbs sampling to estimate marginals
void gibbsSamplingForMarginalEstimation(const GraphicalModel& gm) {
  MarginalVisitor visitor(gm);

  // extend the visitor to sample first order marginals
  for(size_t j = 0; j < gm.numberOfVariables(); ++j) {
    visitor.addMarginal(j);
  }

  // extend the visitor to sample certain second order marginals
  for(size_t j = 0; j < gm.numberOfVariables() - 1; ++j) {
    size_t variableIndices[] = {j, j + 1};
    visitor.addMarginal(variableIndices, variableIndices + 2);
  }

  // sample
  Gibbs gibbs(gm);
  gibbs.infer(visitor);

  // output sampled first order marginals
  std::cout << "sampled_first_order_marginals:" << std::endl;
  for(size_t j = 0; j < gm.numberOfVariables(); ++j) {
    std::cout << "x" << j << ":\n";
    for(size_t k = 0; k < 2; ++k) {
      const double p = static_cast<double>(visitor.marginal(j)(k)) / visitor.numberOfSamples();
      std::cout << p << '\n';
    }
    std::cout << std::endl;
  }

  // output sampled second order marginals
  std::cout << "sampled_second_order_marginals:" << std::endl;
  for(size_t j = gm.numberOfVariables(); j < visitor.numberOfMarginals(); ++j) {
    std::cout << "x" << visitor.marginal(j).variableIndex(0)
      << ",x" << visitor.marginal(j).variableIndex(1)
      << ":\n";
    for(size_t x = 0; x < 2; ++x)
      for(size_t y = 0; y < 2; ++y) {
        const double p = static_cast<double>(visitor.marginal(j)(x, y)) / visitor.numberOfSamples();
        std::cout << p << '\n';
      }
    std::cout << std::endl;
  }
}

int main() {

```

```

    GraphicalModel gm;
    buildGraphicalModel(gm);
    gibbsSamplingForOptimization(gm);
    gibbsSamplingForMarginalEstimation(gm);

    return 0;
}

```

### 3.12.2 Swendsen-Wang

Swendsen-Wang sampling and its generalizations [3] are instances of the Metropolis-Hastings algorithm, with a proposal distribution that allows the labels of multiple variables to change in a single step. Usage of the class template `SwendsenWang` is similar to that of `Gibbs` and is shown in a separate example below.

The implementation in OpenGM is based on [3]. It considers the adjacency graph of variables in which two variables are connected if and only if (i) both variables can attain the same number of labels and (ii) there exists at least one factor that depends on at least both variables. For each such pair of variables, we estimate the joint marginal by accumulating all factors that depend on at least one of the variables, a procedure that is not practical for all graphical models. From this estimate of the marginal, we compute the probability of the two variables taking on equal and unequal labels, respectively.

In each iteration, we consider the adjacency graph of variables above and remove all edges between variables that are currently assigned different labels. From the remaining connected components, edges are removed at random, according to the probabilities estimated before. Finally, one of the now possibly even smaller connected components and a proposed new label are drawn uniformly at random, and this relabeling is subject to the Metropolis-Hastings acceptance procedure. The C++ implementation is more efficient and does not involve dynamic manipulations of graphs.

```

#include <vector>
#include <iostream>

#include <opengm/graphicalmodel/space/simplediscretespace.hxx>
#include <opengm/functions/explicit.function.hxx>
#include <opengm/functions/potts.hxx>
#include <opengm/operations/multiplier.hxx>
#include <opengm/operations/maximizer.hxx>
#include <opengm/graphicalmodel/graphicalmodel.hxx>
#include <opengm/inference/swendsenwang.hxx>

typedef opengm::SimpleDiscreteSpace<> Space;
typedef opengm::ExplicitFunction<double> ExplicitFunction;
typedef opengm::PottsFunction<double> PottsFunction;
typedef opengm::meta::TypeListGenerator<ExplicitFunction, PottsFunction>::type FunctionTypes;
typedef opengm::GraphicalModel<double, opengm::Multiplier, FunctionTypes, Space> GraphicalModel;
typedef opengm::SwendsenWang<GraphicalModel, opengm::Maximizer> SwendsenWang;
typedef opengm::SwendsenWangMarginalVisitor<SwendsenWang> MarginalVisitor;

// build a Markov Chain with 10 binary variables in which
// - the first variable is more likely to be labeled 1 than 0
// - neighboring variables are more likely to have similar labels than dissimilar
void buildGraphicalModel(GraphicalModel& gm) {
    const size_t numberOfVariables = 10;
    const size_t numberOfLabels = 2;
    Space space(numberOfVariables, numberOfLabels);
    gm = GraphicalModel(space);

    // add 1st order function
    GraphicalModel::FunctionIdentifier fid1;
}

```

```

    ExplicitFunction f(&numberOfLabels, &numberOfLabels + 1);           32
    f(0) = 0.2;                                                         33
    f(1) = 0.8;                                                         34
    fid1 = gm.addFunction(f);                                           35
}                                                                         36
                                                                           37
// add 2nd order function                                             38
GraphicalModel::FunctionIdentifier fid2;                               39
{                                                                         40
    const double probEqual = 0.7;                                       41
    const double probUnequal = 0.3;                                       42
    PottsFunction f(2, 2, probEqual, probUnequal);                       43
    fid2 = gm.addFunction(f);                                           44
}                                                                         45
                                                                           46
// add 1st order factor (at first variable)                            47
{                                                                         48
    size_t variableIndices[] = {0};                                       49
    gm.addFactor(fid1, variableIndices, variableIndices + 1);           50
}                                                                         51
                                                                           52
// add 2nd order factors                                             53
for(size_t j = 0; j < numberOfVariables - 1; ++j) {                   54
    size_t variableIndices[] = {j, j+1};                                   55
    gm.addFactor(fid2, variableIndices, variableIndices + 2);           56
}                                                                         57
}                                                                         58
                                                                           59
// use SwendsenWang sampling for the purpose of finding the most probable labeling
void swendsenWangSamplingForOptimization(const GraphicalModel& gm) {     60
    const size_t numberOfSamplingSteps = 1e4;                             61
    const size_t numberOfBurnInSteps = 1e4;                               62
    SwendsenWang::Parameter parameter(numberOfSamplingSteps, numberOfBurnInSteps); 63
    SwendsenWang swendsenWang(gm, parameter);                           64
    swendsenWang.infer();                                               65
    std::vector<size_t> argmax;                                           66
    swendsenWang.arg(argmax);                                           67
                                                                           68
    std::cout << "most_probable_labeling_sampled:\n";                    69
    for(size_t j = 0; j < argmax.size(); ++j) {                           70
        std::cout << argmax[j] << ", ";                                  71
    }                                                                     72
    std::cout << "\b\b)" << std::endl;                                    73
}                                                                         74
                                                                           75
// use SwendsenWang sampling to estimate marginals                    76
void swendsenWangSamplingForMarginalEstimation(const GraphicalModel& gm) { 77
    MarginalVisitor visitor(gm);                                         78
                                                                           79

    // extend the visitor to sample first order marginals              80
    for(size_t j = 0; j < gm.numberOfVariables(); ++j) {                 81
        visitor.addMarginal(j);                                          82
    }                                                                     83
                                                                           84

    // extend the visitor to sample certain second order marginals     85
    for(size_t j = 0; j < gm.numberOfVariables() - 1; ++j) {            86
        size_t variableIndices[] = {j, j + 1};                           87
        visitor.addMarginal(variableIndices, variableIndices + 2);       88
    }                                                                     89
                                                                           90

    // sample                                                            91
    SwendsenWang swendsenWang(gm);                                       92
    swendsenWang.infer(visitor);                                         93
                                                                           94

    // output sampled first order marginals                              95
    std::cout << "sampled_first_order_marginals:" << std::endl;         96
    for(size_t j = 0; j < gm.numberOfVariables(); ++j) {                 97

```

```

std::cout << "x" << j << ": ";
for(size_t k = 0; k < 2; ++k) {
    const double p = static_cast<double>(visitor.marginal(j)(k)) / visitor.numberOfSamples();
    std::cout << p << " ";
}
std::cout << std::endl;
}
// output sampled second order marginals
std::cout << "sampled_second_order_marginals:" << std::endl;
for(size_t j = gm.numberOfVariables(); j < visitor.numberOfMarginals(); ++j) {
    std::cout << "x" << visitor.marginal(j).variableIndex(0)
        << ",x" << visitor.marginal(j).variableIndex(1)
        << ": ";
    for(size_t x = 0; x < 2; ++x)
        for(size_t y = 0; y < 2; ++y) {
            const double p = static_cast<double>(visitor.marginal(j)(x, y)) / visitor.numberOfSamples();
            std::cout << p << " ";
        }
    std::cout << std::endl;
}
}
}
int main() {
    GraphicalModel gm;

    buildGraphicalModel(gm);
    swendsenWangSamplingForOptimization(gm);
    swendsenWangSamplingForMarginalEstimation(gm);

    return 0;
}

```

### 3.13 Brute Force Search

OpenGM provides an implementation of brute force exhaustive search, mainly for the purpose of testing other algorithms and to solve small problems with minimal overhead. In fact, the header file `include/opengm/inference/bruteforce.hxx` is an excellent starting point for implementing custom algorithms (Section 6.4).

```

#include <stdlib.h>
#include <opengm/graphicalmodel/graphicalmodel.hxx>
#include <opengm/operations/adder.hxx>
#include <opengm/operations/minimizer.hxx>
#include <opengm/inference/bruteforce.hxx>
typedef opengm::GraphicalModel<double, opengm::Adder> GraphicalModelType;
typedef opengm::Bruteforce<GraphicalModelType,opengm::Minimizer> Bruteforce;
gm;// ...Bruteforce bf(gm);bf.infer();
std::cout << "value:" << bf.value() << std::endl;

```

### 3.14 Wrappers around Other Libraries

#### 3.14.1 libDAI\*

OpenGM wrappers (Tab. 3.2) around the inference algorithms of libDAI [28, 29] are collected in the namespace `opengm::external::libdai`: These classes derive from `opengm::external::libdai::Inference` which itself derives from `opengm::Inference`. Therefore, wrapped algorithms can be used exactly like native OpenGM algorithms:



Algorithm	Wrapper class
Belief Propagation	Bp
TRBP	TreeReweightedBp
Double Loop Generalized BP	DoubleLoopGeneralizedBP
Fractional BP	FractionalBp
Loop Corrected BP	LoopCorrectedBp
Tree Expectation Propagation	TreeExpectationPropagation
Exact Inference	Exact
Junction Tree	JunctionTree
Gibbs Sampling	Gibbs
Mean Field	MeanField

Table 3.2: OpenGM classes wrapping libDAI inference algorithms.

### 3.14.1.1 Belief Propagation

```

#include <opengm/inference/external/libdai/bp.hxx> 1
#include <opengm/operations/maximizer.hxx> 2
// set up the optimizer (bp) 3
typedef external::libdai::Bp<Model, Maximizer> Bp; 4
5
const size_t maxIterations=100; 6
const double damping=0.0; 7
const double tolerance=0.000001; 8
// Bp::UpdateRule = PARALL | SEQFIX | SEQRND | SEQMAX 9
Bp::UpdateRule updateRule= PARALL; 10
const size_t verbose=0; 11
size_t verboseLevel=0; 12
Bp::Parameter parameter(maxIterations, damping,tolerance,updateRule,verboseLevel); 13
Bp bp(gm, parameter); 14
// optimize (approximately) 15
bp.infer( ); 16
// obtain the (approximate) argmax 17
vector<size_t> labeling(numberOfVariables); 18
bp.arg(labeling); 19

```

### 3.14.1.2 TRBP

```

#include <opengm/inference/external/libdai/tree_reweighted_bp.hxx> 1
#include <opengm/operations/maximizer.hxx> 2
// set up the optimizer (trbp) 3
typedef external::libdai::TreeReweightedBp<Model, Maximizer> Trbp; 4
5
const size_t maxIterations=100; 6
const double damping=0.0; 7
const double tolerance=0.000001; 8
const size_t ntrees=10; 9
// Trbp::UpdateRule = PARALL | SEQFIX | SEQRND | SEQMAX 10
Trbp::UpdateRule updateRule= PARALL; 11
const size_t verboseLevel=0; 12
size_t verboseLevel=0; 13
Trbp::Parameter parameter(maxIterations, damping,tolerance,ntrees,updateRule,verboseLevel); 14
Trbp trbp(gm, parameter); 15
// optimize (approximately) 16
trbp.infer( ); 17
// obtain the (approximate) argmax 18
vector<size_t> labeling(numberOfVariables); 19
trbp.arg(labeling); 20

```

## 3.14.1.3 Double Loop Generalized BP

```

#include <opengm/inference/external/libdai/fractional_bp.hxx> 1
#include <opengm/operations/minimizer.hxx> 2
// set up the optimizer (fractional bp) 3
typedef external::libdai::DoubleLoopGeneralizedBP<Model, Minimizer> DoubleLoopGeneralizedBP; 4
5
const bool doubleloop=1; 6
// DoubleLoopGeneralizedBP::Clusters=MIN | BETHE | DELTA | LOOP 7
const DoubleLoopGeneralizedBP::Clusters clusters=BETHE; 8
const size_t loopdepth = 3; 9
// DoubleLoopGeneralizedBP::Init = UNIFORM | RANDOM 10
const DoubleLoopGeneralizedBP::Init init=UNIFORM; 11
const size_t maxiter=10000; 12
const double tolerance=1e-9; 13
const size_t verboseLevel=0; 14
DoubleLoopGeneralizedBP::Parameter parameter( loopdepth, init, maxiter, tolerance,verboseLevel); 15
DoubleLoopGeneralizedBP gdlbp(gm, parameter); 16
// optimize (approximately) 17
gdlbp.infer( ); 18
// obtain the (approximate) argmin 19
vector<size_t> labeling(numberOfVariables); 20
gdlbp.arg(labeling); 21

```

## 3.14.1.4 Fractional BP

```

#include <opengm/inference/external/libdai/fractional_bp.hxx> 1
#include <opengm/operations/maximizer.hxx> 2
// set up the optimizer (fractional bp) 3
typedef external::libdai::FractionalBp<Model, Maximizer> FractionalBp; 4
5
const size_t maxIterations=100; 6
const double damping=0.0; 7
const double tolerance=0.000001; 8
// FractionalBp::UpdateRule = PARALL | SEQFIX | SEQRND | SEQMAX 9
FractionalBp::UpdateRule updateRule= PARALL; 10
const size_t verboseLevel=0; 11
size_t verboseLevel=0; 12
FractionalBp::Parameter parameter(maxIterations, damping,tolerance,updateRule,verboseLevel); 13
FractionalBp fbp(gm, parameter); 14
// optimize (approximately) 15
fbp.infer( ); 16
// obtain the (approximate) argmax 17
vector<size_t> labeling(numberOfVariables); 18
fbp.arg(labeling); 19

```

## 3.14.1.5 TreeExpectationPropagation

```

#include <opengm/inference/external/libdai/fractional_bp.hxx> 1
#include <opengm/operations/maximizer.hxx> 2
// set up the optimizer (fractional bp) 3
typedef external::libdai::TreeExpectationPropagation<Model, Maximizer> TreeExpectationPropagation; 4
5
//TreeExpectationPropagation::TreeEpType = ORG | ALT 6
TreeExpectationPropagation::TreeEpType treeEpTyp=ORG; 7
const size_t maxiter=10000, 8
const double tolerance=1e-9, 9
size_t verboseLevel=0 10
TreeExpectationPropagation::Parameter parameter(treeEpTyp,maxiter,tolerance,verboseLevel); 11
TreeExpectationPropagation treeep(gm, parameter); 12

```

```

// optimize (approximately) 13
treep.infer( ); 14
// obtain the (approximate) argmax 15
vector<size_t> labeling(numberOfVariables); 16
treep.arg(labeling); 17

```

### 3.14.1.6 Exact

```

#include <opengm/inference/external/libdai/exact.hxx> 1
#include <opengm/operations/minimizer.hxx> 2
// set up the optimizer (junction tree) 3
typedef external::libdai::Exact<Model, Minimizer> Exact; 4
5
size_t verboseLevel=0; 6
Exact::Parameter parameter(verboseLevel); 7
Exact exact(gm, parameter); 8
// optimize (to global optimum) 9
exact.infer( ); 10
// obtain the argmin 11
vector<size_t> labeling(numberOfVariables); 12
exact.arg(labeling); 13

```

### 3.14.1.7 Junction Tree

```

#include <opengm/inference/external/libdai/junction_tree.hxx> 1
#include <opengm/operations/minimizer.hxx> 2
// set up the optimizer (exact) 3
typedef external::libdai::JunctionTree<Model, Minimizer> JunctionTree; 4
5
// JunctionTree::UpdateRule = HUGIN | SHSH 6
JunctionTree::UpdateRule updateRule=HUGIN; 7
// JunctionTree::Heuristic = MINFILL | WEIGHTEDMINFILL | MINNEIGHBORS 8
JunctionTree::Heuristic heuristic=MINFILL; 9
size_t verboseLevel=0; 10
JunctionTree::Parameter parameter(updateRule, heuristic, verboseLevel); 11
JunctionTree jt(gm, parameter); 12
// optimize (to global optimum) 13
jt.infer( ); 14
// obtain the argmin 15
vector<size_t> labeling(numberOfVariables); 16
jt.arg(labeling); 17

```

### 3.14.1.8 Gibbs

```

#include <opengm/inference/external/libdai/gibbs.hxx> 1
#include <opengm/operations/minimizer.hxx> 2
// set up the optimizer (gibbs sampler) 3
typedef external::libdai::Gibbs<Model, Minimizer> Gibbs; 4
5
const size_t maxiter=10000; 6
const size_t burnin=100; 7
const size_t restart_=10000; 8
const size_t verbose=0; 9
Gibbs::Parameter parameter(maxiter, burnin, restart_, verboseLevel); 10
Gibbs gibbs(gm, parameter); 11
// optimize 12
gibbs.infer( ); 13
// obtain the argmin 14

```

```
vector<size_t> labeling(numberOfVariables);           15
gibbs.arg(labeling);                                16
```

### 3.14.1.9 Mean Field

```
#include <opengm/inference/external/libdai/mean_field.hxx>           1
#include <opengm/operations/minimizer.hxx>                           2
// set up the optimizer (mean field)                                   3
typedef external::libdai::MeanField<Model, Minimizer> MeanField;    4
                                                                    5
const size_t maxiter=10000;                                          6
const double damping=0.2;                                           7
const double tolerance=1e-9;                                        8
// MeanField::UpdateRule = NAIVE | HARDSPIN                          9
const MeanField::UpdateRule updateRule= NAIVE;                     10
// MeanField::Init = UNIFORM | RANDOM                                11
const MeanField::Init init=UNIFORM;                                 12
const size_t verboseLevel=0;                                        13
MeanField::Parameter parameter(maxiter,damping,tolerance,updateRule,verboseLevel); 14
MeanField mf(gm, parameter);                                        15
// optimize (approximately)                                         16
mf.infer( );                                                       17
// obtain the (approximate) argmin                                   18
vector<size_t> labeling(numberOfVariables);                          19
mf.arg(labeling);                                                  20
```

### 3.14.2 MRF-LIB\*

MRF-LIB [34, 33] is a highly optimized inference library. It is restricted to energy minimization problem and second-order grid models with a four neighborhood. Many important computer vision problems belong to this class. MRF-LIB implements Loopy Belief Propagation (LBP), Sequential Belief Propagation (BP-S), Tree Re-weighted Belief Propagation (TRW-S), Iterated Conditional Modes (ICM),  $\alpha$ -Expansion, and  $\alpha\beta$ -Swap.

OpenGM makes all these implementations available in two alternative ways: On the one hand, a graphical model constructed in OpenGM copied into the native MRF-LIB data-structure for which highly optimized code exists. On the other hand, a view function can be used internally that makes OpenGM models look to the MRF-LIB code as if they were stored in the native data structure. No data is copied in this case. In fact, MRF-LIB can even be beneficial from the OpenGM data structures that distinguishes between functions and factors, e.g. for the photo montage models of the MRF-benchmark [34]. Here, the OpenGM view is more efficient because the repeated use of the same function with multiple factors reduces cache misses.

The access to MRF-LIB is via a the single class template `opengm::external::MRFLIB`. Important parameters are:

- The inference algorithm `inferenceType_`. Possible choices are ICM, EXPANSION, SWAP, MAX-PRODBP, TRWS, and BPS.
- The type of energy function `energyType_`. Choices are an OpenGM view VIEW, a single table for each factor TABLES, truncated linear label differences TL1, truncated quadratic label differences TL2, or a single table weighted independently for each pairwise factor WEIGHT-EDTABLE.
- Maximum number of iterations `numberOfIterations_`.
- For TRW-S, OpenGM stops if `fabs(value - bound) / max(fabs(value), 1) < trwsTolerance_`.

## 3.14.2.1 ICM

```

#include <stdlib.h> 1
#include <opengm/graphicalmodel/graphicalmodel.hxx> 2
#include <opengm/operations/adder.hxx> 3
#include <opengm/operations/minimizer.hxx> 4
#include <opengm/inference/external/mrflib.hxx> 5
6
typedef opengm::GraphicalModel<double, opengm::Adder> GraphicalModelType; 7
typedef opengm::external::MRFLIB<GraphicalModelType> MRFLIB; 8
9
GraphicalModelType gm; 10
// ... 11
MRFLIB::Parameter para; 12
para.inferenceType_ = MRFLIB::Parameter::ICM; 13
para.energyType_ = MRFLIB::Parameter::WEIGHTEDTABLE; 14
para.numberOfIterations_ = 10; 15
16
MRFLIB mrf(gm,para); 17
mrf.infer(); 18
19
std::cout << "value:_" << mrf.value() << std::endl; 20

```

3.14.2.2  $\alpha$ -Expansion

```

#include <stdlib.h> 1
#include <opengm/graphicalmodel/graphicalmodel.hxx> 2
#include <opengm/operations/adder.hxx> 3
#include <opengm/operations/minimizer.hxx> 4
#include <opengm/inference/external/mrflib.hxx> 5
6
typedef opengm::GraphicalModel<double, opengm::Adder> GraphicalModelType; 7
typedef opengm::external::MRFLIB<GraphicalModelType> MRFLIB; 8
9
GraphicalModelType gm; 10
// ... 11
MRFLIB::Parameter para; 12
para.inferenceType_ = MRFLIB::Parameter::EXPANSION; 13
para.energyType_ = MRFLIB::Parameter::TL1; 14
para.numberOfIterations_ = 10; 15
16
MRFLIB mrf(gm,para); 17
mrf.infer(); 18
19
std::cout << "value:_" << mrf.value() << std::endl; 20

```

3.14.2.3  $\alpha\beta$ -Swap

```

#include <stdlib.h> 1
#include <opengm/graphicalmodel/graphicalmodel.hxx> 2
#include <opengm/operations/adder.hxx> 3
#include <opengm/operations/minimizer.hxx> 4
#include <opengm/inference/external/mrflib.hxx> 5
6
typedef opengm::GraphicalModel<double, opengm::Adder> GraphicalModelType; 7
typedef opengm::external::MRFLIB<GraphicalModelType> MRFLIB; 8
9
GraphicalModelType gm; 10
// ... 11
MRFLIB::Parameter para; 12
para.inferenceType_ = MRFLIB::Parameter::SWAP; 13
para.energyType_ = MRFLIB::Parameter::TL1; 14
para.numberOfIterations_ = 10; 15
16

```

```

MRFLIB mrf(gm,para);
mrf.infer();

std::cout << "value:_" << mrf.value() << std::endl;

```

#### 3.14.2.4 LBP

```

#include <stdlib.h>
#include <opengm/graphicalmodel/graphicalmodel.hxx>
#include <opengm/operations/adder.hxx>
#include <opengm/operations/minimizer.hxx>
#include <opengm/inference/external/mrflib.hxx>

typedef opengm::GraphicalModel<double, opengm::Adder> GraphicalModelType;
typedef opengm::external::MRFLIB<GraphicalModelType> MRFLIB;

GraphicalModelType gm;
// ...
MRFLIB::Parameter para;
para.inferenceType_ = MRFLIB::Parameter::MAXPRODBP;
para.energyType_ = MRFLIB::Parameter::TABLES;
para.numberOfIterations_ = 10;

MRFLIB mrf(gm,para);
mrf.infer();

std::cout << "value:_" << mrf.value() << std::endl;

```

#### 3.14.2.5 BP-S

```

#include <stdlib.h>
#include <opengm/graphicalmodel/graphicalmodel.hxx>
#include <opengm/operations/adder.hxx>
#include <opengm/operations/minimizer.hxx>
#include <opengm/inference/external/mrflib.hxx>

typedef opengm::GraphicalModel<double, opengm::Adder> GraphicalModelType;
typedef opengm::external::MRFLIB<GraphicalModelType> MRFLIB;

GraphicalModelType gm;
// ...
MRFLIB::Parameter para;
para.inferenceType_ = MRFLIB::Parameter::BPS;
para.energyType_ = MRFLIB::Parameter::VIEW;
para.numberOfIterations_ = 10;

MRFLIB mrf(gm,para);
mrf.infer();

std::cout << "value:_" << mrf.value() << std::endl;

```

#### 3.14.2.6 TRW-S

```

#include <stdlib.h>
#include <opengm/graphicalmodel/graphicalmodel.hxx>
#include <opengm/operations/adder.hxx>
#include <opengm/operations/minimizer.hxx>
#include <opengm/inference/external/mrflib.hxx>

typedef opengm::GraphicalModel<double, opengm::Adder> GraphicalModelType;
typedef opengm::external::MRFLIB<GraphicalModelType> MRFLIB;

```

```
GraphicalModelType gm; 10
// ... 11
MRFLIB::Parameter para; 12
para.inferenceType_ = MRFLIB::Parameter::TRWS; 13
para.energyType_ = MRFLIB::Parameter::VIEW; 14
para.numberofIterations_ = 10; 15
para.trwsTolerance_ = 1 16
17
MRFLIB mrf(gm,para); 18
mrf.infer(); 19
20
std::cout << "value:_" << mrf.value() << "_bound:_" << mrf.bound() << std::endl; 21
```





# Chapter 4

## Command Line Tools

### 4.1 Performing Inference with Command Line Tools

The OpenGM command line interface provides a simple way to perform inference in the min-sum and max-product semi-ring. This interface is split up into several tools, each corresponding to a graphical model type and suitable inference algorithms<sup>1</sup>.

As described also in Section 2.2, the type of a graphical model is defined by: (i) *a value type*, (ii) *an operator type*, (iii) *a set of function types*, and (iv) *a space type*.

The type of an inference algorithm is defined by: (i) *a graphical model type*, and (ii) *an accumulation type* (Chapter 3).

An command line tool can be built for each reasonable combination. Since the set of possible combinations is large, the original source includes only a few examples. Is is discussed in Section 4.2 how one can add more combinations. The set of function types is set to the same default function types which are supported even by the most basic OpenGM build: (i) *ExplicitFunction*, (ii) *PottsFunction*, (iii) *PottsNFunction*, (iv) *PottsGFunction*, (v) *TruncatedSquaredDifferenceFunction*, (vi) *TruncatedAbsoluteDifferenceFunction*. Adding and removing functions is easy (Section 4.2). The



Figure 4.1: To use the opengm command line tool, one should proceed in the following: First generate the model with MatLab, in C, or any other way, than use the command line tool to generate a protocol file. This can be viewed by hdfview or loaded into MatLab, C, or the programming language of your choice.

basic work-flow for using the command line tool is sketched in Fig. 4.1. After a graphical model is generated in C++ or MATLAB, a command line tool can be used to perform inference. The command line tool can protocol the inference procedure and save optimal configurations, bounds and execution times and dump these into an HDF5 file. This file can be inspected using hdfview, read by MATLAB and C++ code or any other program that supports HDF5. Parameters of the command line tool specify parameters of inference algorithms. Building a graphical model, writing it to an HDF5 file and then performing inference by means of a command line tool can be an appealing workflow for experimental research. It makes it easy to distribute inference tasks

---

<sup>1</sup>Note that a single executable would become inconveniently large due to the various combinations of graphical models and inference algorithms

for multiple graphical models to the nodes of a cluster computer and encourages the exchange of graphical models.

**NOTE:** Command line tools should be built in release mode (set `CMAKE_BUILD_TYPE` to `RELEASE` and compile again) to run at peak performance.

### 4.1.1 Basic Usage

No matter which executable is chosen to perform inference, the command line syntax is always the same. A good starting point is the parameter `-h` or `--help` which can be used to show top level options:

```
> opengm_min_sum -h
This is the help for the commandline interface of the opengm library
Usage: opengm [arguments]
arguments:
  short name  long name          needed  description
  -h          --help           no      used to activate help output
  -v          --verbose        no      used to activate verbose output
  --modelinfo --modelinfo    no      used to print detailed informations about the
  specified model
  -m          --model         yes     used to specify the desired model. Usage: filename
  :dataset
  -a          --algorithm     no      used to specify the desired algorithm
  permitted values: GRAPHCUT; ALPHAEXPANSION;
  ALPHABETASWAP; QPBO; LPCPLEX; TRWS; MRF;
  BELIEFPROPAGATION; TREP; ASTAR; LAZYFLIPPER;
  -o          --outputfile    no      used to specify the desired outputfile for the
  computed results
  default value: PRINT ON SCREEN
  -e          --evaluate      no      used to evaluate the selected model with the
  specified input file
```

Some options, when selected, enable more specific options. These need to be set after the top level option from which they must be separated by a space. These are all top level parameters:

- h \ --help** Prints the help on screen. If used in combination with parameter `-a` the help for the corresponding algorithm will be printed.
- v \ --verbose** If used in combination with parameter `-p` (a standard parameter for every algorithm) verbose information on every iteration will be printed on screen (see 4.1.2 for details).
- modelinfo** Prints some information about a model specified with parameter `-m`.
- m \ --model** Used to load a model. Use `filename:dataset` to load a model from the dataset of a given hdf5 file.
- a \ --algorithm** Specify the algorithm which will be used to perform inference on a given model.
- o \ --outputfile** Specify a name for the file in which the computed results will be stored. Depending on the file extension the results will be stored in a text file with comma separated values (`.txt`) or in a hdf5 file (`.h5`).
- e \ --evaluate** Evaluates a model specified with the parameter `-m` with the vector stored in `filename:dataset`.

Each inference algorithm provides its own set of control parameters. These can be looked up by using the option `-h` in combination with the option `-a`, e.g.

```
> opengm_min_sum -h -a ALPHAEXPANSION
Printing Help for inference caller ALPHAEXPANSION
Description:
detailed description of AlphaExpansion caller...
arguments:
  short name  long name          needed  description
```

```

-p          --protocolate      no      used to enable protocolation mode. Usage: "-p N"
          where every Nth iteration step will be
          protocolated. If N = 0 only the final results
          will be protocolated.
          --scale              no      default value: 0
          Add description for scale here!!!!.
          default value: 1
-mf         --maxflow          no      Add description for MinSTCut here!!!!.
          permitted values: KOLMOGOROV; BOOST_KOLMOGOROV;
          BOOST_PUSH_RELABEL; BOOST_EDMONDS_KARP;
          default value: KOLMOGOROV
          --maxIt              no      Maximum number of iterations.
          default value: 1000
          --labelInitialType  no      select the desired initial label
          permitted values: DEFAULT; RANDOM; LOCALOPT;
          EXPLICIT;
          default value: DEFAULT
          --orderType         no      select the desired order
          permitted values: DEFAULT; RANDOM; EXPLICIT;
          default value: DEFAULT
          --randSeedOrder     no      Add description for randSeedOrder here!!!!.
          default value: 0
          --randSeedLabel     no      Add description for randSeedLabel here!!!!.
          default value: 0
          --labelorder        no      location of the file containing a vector which
          specifies the desired label order
          --label              no      location of the file containing a vector which
          specifies the desired label

```

As an example, we call the command line tool `opengm_min_sum`, the  $\alpha$ -expansion inference algorithm for a model stored in the file `MyModel.h5`, in the HDF5 dataset `gm1`:

```
> opengm_min_sum -m MyModel.h5:gm1 -a ALPHAEXPANSION
```

The inference algorithms shipped with OpenGM have a complete set of default parameters. Thus, all parameters are optional. Each algorithm can be called by simply passing its name and a model to a command line tool. The order in which parameters are provided is irrelevant. A more complex example can be found in Section 4.1.3.

### 4.1.2 Keeping a Protocol

A special parameter is the parameter `-p` or `--protocolate` which is available for all inference algorithms. If this parameter is set to a value  $N > 0$  then protocolation mode will be enabled which means that for every  $N$ th iteration the current value, bound and passed time will be stored in addition to the computed optimal state. If the verbose parameter is used in combination with the `protocolate` parameter then the values for every  $N$ th iteration will also be printed as console output.

### 4.1.3 Example for Usage of the Command Line Tool

We have converted models of the MRF-Benchmark (<http://vision.middlebury.edu/MRF/results/>) to the OpenGM model format and now want to use the `tsukuba` model to show some of the capabilities of the command line tools. The model is stored in the file `tsu-gm.h5` in the dataset `gm`.

First of all we are interested in some basic information about our model, so we run

```
> opengm_min_sum -m tsu-gm.h5:gm --modelinfo
```

which produces the following output:

```
loading model: tsu-gm.h5
using dataset: gm
Model Info
```

```

-----
Number of Variables   : 110592
Number of States     : 16
Model-Order          : 2
Number of Factors    : 331104
Number of Function Types : 6
Acyclic Model        : 0
Chain Model          : 0
Grid Model           : 1

```

This tells us that our model consist of 110592 variables with a number of 16 states for each variable. The maximum order of a factor is 2 with a total number of 331104 factors. The model type includes 6 function types. Our model is not acyclic, has not a chain structure but a grid structure.

In the second place we want to run our model with the mrf inference algorithm to compare our computed results with the results presented on the MRF-Benchmark website. Therefore we run the following command:

```
> opengm_min_sum -a MRF -m tsu-gm.h5:gm --inference TRWS --energy TL1 -o tsu-gm-results.txt -p 1 -v --maxIt 5
```

Thus the inference algorithm mrf is called with our model and uses the implementation of the trws algorithm with the energy form truncated linear. Furthermore we select the file tsu-gm-results.h5 to store the computed results and protocolate each iteration step. To not blast the limits of this document we set a maximum iteration number of 5. The verbose option is set to also print all computed results on screen and the maximum number of iterations is set to 5 to not blow up the output. This leads to the following output:

```

loading model: tsu-gm.h5
using dataset: gm
running MRF caller
T: 2
Begin :      0 Value : 5.046306e+06 Bound : 1.546425e-321
Step  :      0 Value : 4.306390e+05 Bound : 3.322034e+05
Step  :      1 Value : 4.075980e+05 Bound : 3.457036e+05
Step  :      2 Value : 3.996550e+05 Bound : 3.514903e+05
Step  :      3 Value : 3.958940e+05 Bound : 3.548906e+05
Step  :      4 Value : 3.934340e+05 Bound : 3.571302e+05
End   :      5 Value : 3.934340e+05 Bound : 3.571302e+05
storing values in file: tsu-gm-results.h5
storing bounds in file: tsu-gm-results.h5
storing times in file: tsu-gm-results.h5
storing corresponding iterations in file: tsu-gm-results.h5
storing optimal states in file: tsu-gm-results.h5

```

The output shows some IO information, and more interesting that our energy function is truncated at 2 and the values and bounds computed at each iteration step these correspond to the results stored in tsu-gm-results.h5 as shown in Figure 4.2 and are also identical to those values presented on the MRF website.

Finally after we have computed the optimal state we can now check if the optimal state coincident with the computed optimal value. This can be done by using the parameter `--evaluate` which gives the following result:

```
> opengm_min_sum -m tsu-gm.h5:gm --evaluate tsu-gm-results.h5:states
loading model: tsu-gm.h5
using dataset: gm
result: 3.934340e+05
```

This is the same value as the one we got from our inference run above. So our optimal state seems to coincident with our computed optimal value.

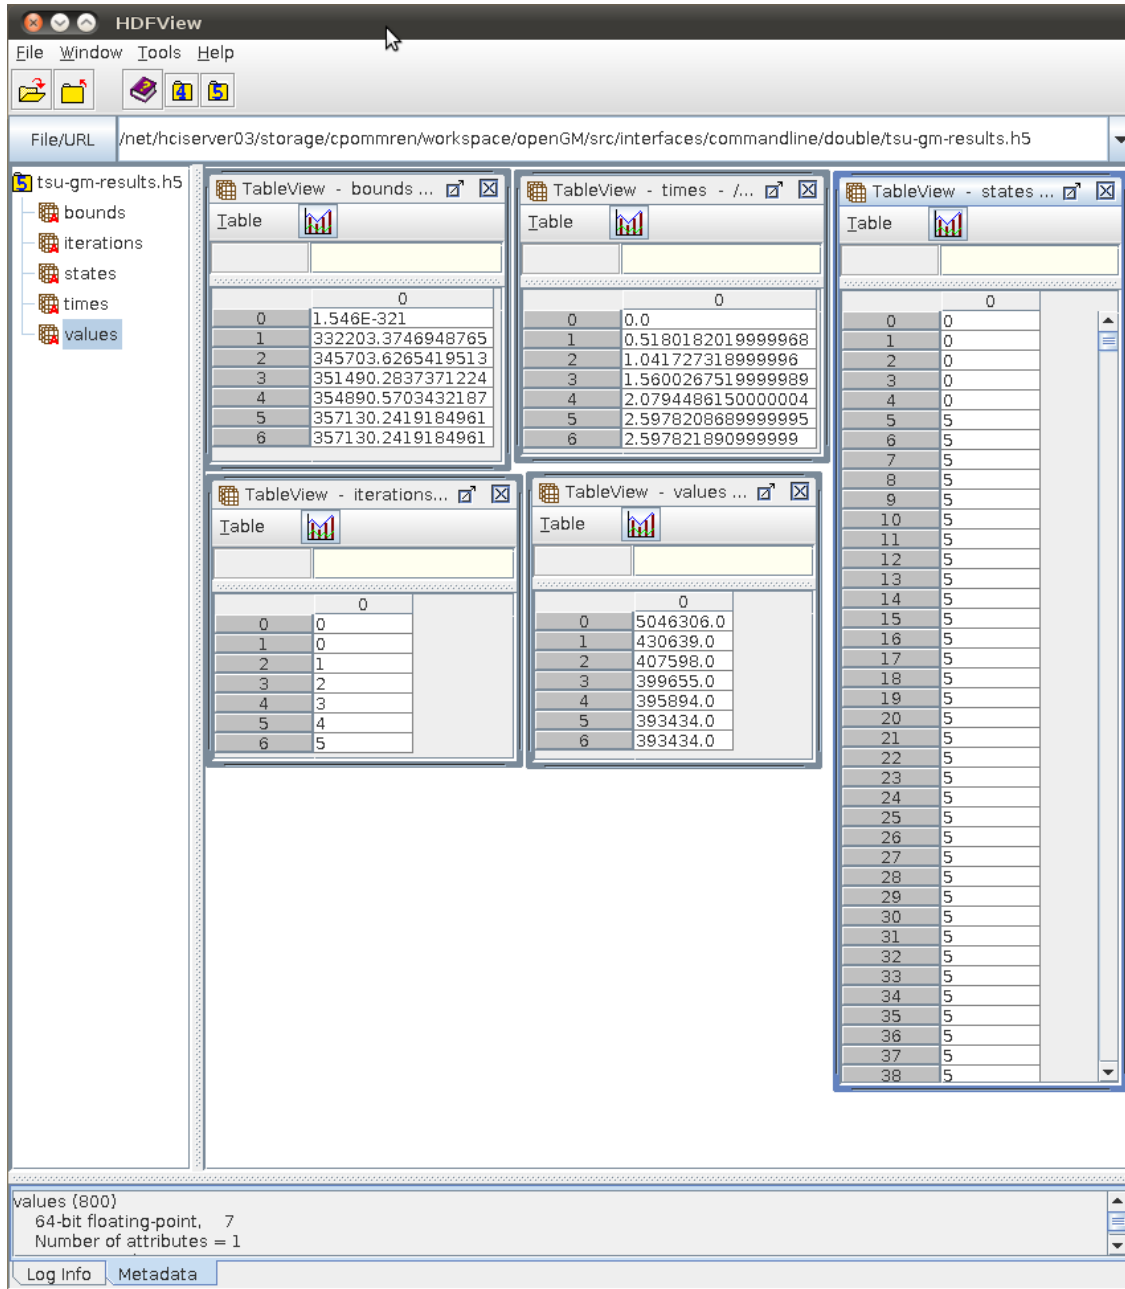


Figure 4.2: Results of the MRF-Benchmark tsukuba computed with OpenGM command line tools with the call `opengm_min_sum -a MRF -m tsu-gm.h5:gm -inference TRWS -energy TL1 -o tsu-gm-results.txt -p 1 -v -maxIt 5`

### 4.1.4 Reading Protocol Files in MatLab

Reading protocol files in MatLab is very simple. After the directory including the additional MatLab scripts has been added to the MatLab path by

```
addpath('~/opengm2/matlab')
```

one can use the functions `marray_load` and `marray_save` to load and save multidimensional arrays into HDF5-files, respectively. The following code load vectors of values, bounds and times stored by the command line tool as well as the final state-vector.

```
addpath('~/opengm2/matlab')

resfile = './tsu-gm-results.h5';
times   = marray_load(resfile,'times');
bounds  = marray_load(resfile,'bounds');
values  = marray_load(resfile,'values');

statevector = marray_load('out-opengm_lbp-gm94.h5','states');
```

## 4.2 Changing Model Type

If you want to use the OpenGM command line tools with our own model, functions, or types, you have to create a new executable which supports this model type. The best way to do this is to look at in the OpenGM source folder (`src/interfaces/commandline/`). There are several example codes for different types of models.

For instance the minimalistic code to create an executable which only supports a graphical model with Value Type *float*, Index Type *unsigned int*, Label Type *unsigned char*, Operator Type *opengm::Adder*, and Function Type *opengm::ExplicitFunction* and the bruteforce inference algorithm with Accumulation Type *opengm::Minimizer* is as follows:

```
#include <iostream> 1
2
// opengGM stuff 3
#include <opengm/graphicalmodel/graphicalmodel.hxx> 4
#include <opengm/graphicalmodel/graphicalmodel.hdf5.hxx> 5
#include <opengm/operations/minimizer.hxx> 6
#include <opengm/operations/adder.hxx> 7
#include <opengm/functions/explicit_function.hxx> 8
9
// commandline IO interface 10
#include "opengm/src/interfaces/commandline/cmd_interface.hxx" 11
12
// inference caller 13
#include "opengm/src/interfaces/common/caller/bruteforce_caller.hxx" 14
15
int main(int argc, char** argv) { 16
    if(argc < 2) { 17
        std::cerr << "At least one input argument required" << std::endl; 18
        std::cerr << "try \"-h\" for help" << std::endl; 19
        return 1; 20
    } 21
22
    typedef float ValueType; 23
    typedef unsigned int IndexType; 24
    typedef unsigned char LabelType; 25
    typedef opengm::Adder OperatorType; 26
    typedef opengm::Minimizer AccType; 27
    typedef opengm::interface::IOCMD InterfaceType; 28
    typedef DiscreteSpace<IndexType, LabelType> SpaceType; 29
30
```

```

// Set functions for graphical model
typedef opengm::meta::TypeListGenerator<
  opengm::ExplicitFunction<ValueType>
>::type FunctionTypeList;

typedef opengm::GraphicalModel<
  ValueType,
  OperatorType,
  FunctionTypeList,
  SpaceType
> GmType;

typedef opengm::meta::TypeListGenerator <
  opengm::interface::BruteforceCaller<interface::IOCMD, GmType, AccType>
>::type InferenceTypeList;

interface::CMDInterface<GmType, InferenceTypeList> interface(argc, argv);
interface.parse();

return 0;
}

```

To change the desired OpenGM Model Type only the lines 23-26 and 29-44 have to be touched.

### 4.3 Adding Inference Methods

To add the support for a new inference method to a command line interface two things have to be done:

1. Write a *Caller Class* for the inference method.
2. Add the caller to the inference type list in the command line tool.

#### Write a Caller Class for the Inference Method:

Say you have followed the guide of section 6.4 and implemented a new inference algorithm called *mySolver*. Then the code for the corresponding caller looks as follows:

```

#include <opengm/opengm.hxx>
// implementation of mySolver
#include <opengm/inference/mySolver.hxx>
// caller base class
#include "inference_caller_base.hxx"
// command line arguments
#include "../argument/argument.hxx"

namespace opengm {
namespace interface {

template <class IO, class GM, class ACC>
class mySolverCaller : public InferenceCallerBase<IO, GM, ACC> {
protected:
  using InferenceCallerBase<IO, GM, ACC>::io_;
  using InferenceCallerBase<IO, GM, ACC>::infer;

  typedef typename opengm::mySolver<GM, ACC> mySolver;
  typedef typename mySolver::Parameter mySolverParameter_;
  typedef typename mySolver::TimingVisitorType TimingVisitorType;

```

```

virtual void                                     24
runImpl(GM& model, StringArgument<>& outputfile, const bool verbose); 25
public:                                           26
  const static std::string name_;                 27
  mySolverCaller(IO& ioIn);                       28
};                                                 29
                                                  30
template <class IO, class GM, class ACC>          31
inline mySolverCaller<IO, GM, ACC>::ICMCaller(IO& ioIn) 32
  : InferenceCallerBase<IO, GM, ACC>("MYSOLVER", 33
    "detailed_description_of_mySolver_Parser...", ioIn) { 34
  // add commandline arguments required by mySolver here... 35
}                                                 36
                                                  37
template <class IO, class GM, class ACC>          38
inline void                                     39
mySolverCaller<IO, GM, ACC>::runImpl(GM& model, 40
  StringArgument<>& outputfile, const bool verbose) { 41
  std::cout << "running_mySolver_caller" << std::endl; 42
  43
  // add code for setting up mySolver here...      44
  45
  // run inference and protocolate results        46
  this-> template infer<ICM, TimingVisitorType, 47
    typename mySolverCaller::Parameter>(model, outputfile, verbose, 48
    mySolverParameter.);                          49
}                                                 50
                                                  51
template <class IO, class GM, class ACC>          52
const std::string mySolverCaller<IO, GM, ACC>::name_ = "MYSOLVER"; 53
  54
} // namespace interface                          55
} // namespace opengm                             56

```

The constructor sets up all necessary information for the command line tools. If you want to add command line parameters which might be required by your solver, this is the place to go (see chapter 4.4 for details).

The function *runImpl()* will be called by the command line tools if a user requests inference with your solver. If your solver requires some setup before *mySolver::infer()* is called, this has to be done in the *runImpl()*-function before the call of the base class function *InferenceCallerBase::infer()*. This base class function takes care of the right call of *mySolver::infer()* and handles storing all user requested values in the specified output file. More examples for existing inference algorithms can be found in `/src/interfaces/common/caller/`.

## Add the Caller to the Inference Type List:

After you have implemented the caller, you have to add it to the inference type list of the command line tool which you want to use. The following code shows the example of chapter 4.2 for a minimalistic command line executable. This time extended by our new caller.

```

#include <iostream>                               1
  2
// opengGM stuff                                  3
#include <opengm/graphicalmodel/graphicalmodel.hxx> 4
#include <opengm/graphicalmodel/graphicalmodel_hdf5.hxx> 5
#include <opengm/operations/minimizer.hxx>        6
#include <opengm/operations/adder.hxx>           7
#include <opengm/functions/explicit_function.hxx> 8
  9
// commandline IO interface                      10
#include "opengm/src/interfaces/commandline/cmd_interface.hxx" 11
  12

```



```

// inference caller
#include "opengm/src/interfaces/common/caller/bruteforce_caller.hxx"

// our new written mySolver caller
#include "mysolver_caller.hxx"

int main(int argc, char** argv) {
    if(argc < 2) {
        std::cerr << "At_least_one_input_argument_required" << std::endl;
        std::cerr << "try_\\"-h\" _for_help" << std::endl;
        return 1;
    }

    typedef float ValueType;
    typedef unsigned int IndexType;
    typedef unsigned char LabelType;
    typedef opengm::Adder OperatorType;
    typedef opengm::Minimizer AccType;
    typedef opengm::interface::IOCMD InterfaceType;
    typedef DiscreteSpace<IndexType, LabelType> SpaceType;

    // Set functions for graphical model
    typedef opengm::meta::TypeListGenerator<
        opengm::ExplicitFunction<ValueType>
    >::type FunctionTypeList;

    typedef opengm::GraphicalModel<
        ValueType,
        OperatorType,
        FunctionTypeList,
        SpaceType
    > GmType;

    typedef opengm::meta::TypeListGenerator <
        opengm::interface::BruteforceCaller<interface::IOCMD, GmType, AccType>,
        opengm::interface::mySolverCaller<interface::IOCMD, GmType, AccType>
    >::type InferenceTypeList;

    interface::CMDInterface<GmType, InferenceTypeList> interface(argc, argv);
    interface.parse();

    return 0;
}

```

All callers which should be supported by a command line toll have to be added to the *InferenceTypeList*.

## 4.4 Adding Parameters

Following the example of a basic caller in chapter 4.3 we will now see how to add command line paramters to the caller to allow the user to specify the behaviour of mySolver.

There are two types of paramters which are supported by the OpenGM command line tools.

First, the *flag-parameter* which can be used to turn the corresponding behaviour *on* or *off*. This type of parameter does not require any further user input as it is basically a bool variable set to either true or false and that why called *BoolArgument*. *BoolArgument* has only one constructor with the signature (bool& storageIn, const std::string& shortNameIn, const std::string& longNameIn, const std::string& descriptionIn) which is identical to the first constructor of the list above, except the requiredIn value as it makes no sense to enforce a flag to be set.

Second, the *value-parameter* which can take a user input to set its value. This type of parameter might have a default value, in which case no further user input is required or it can be tagged as required. In this case the user has to enter a value which has to be entered after the parameter and

has to be separated by a space. Additionally a parameter might have a list of permitted values. If so, the execution of the command line interface will be aborted if the user doesn't enter a value which is part of the permitted values.

The following value-parameter types are supported: (i) *FloatArgument*, (ii) *DoubleArgument*, (iii) *IntArgument*, (iv) *Size\_TArgument*, (v) *StringArgument*, (vi) *VectorArgument*.

All arguments except *BoolArgument* have four different constructors which allow to specify the behaviour of the argument.

The first constructor has the signature (`T& storageIn, const std::string& shortNameIn, const std::string& longNameIn, const std::string& descriptionIn, const bool requiredIn = false`). Where `storageIn` is a reference to the variable in which the user input shall be stored, `shortNameIn` is the short cut name of the parameter (if an empty string is passed as `shortNameIn`, the user can only use the long name to set the parameter), `longNameIn` is the long name of the parameter, `descriptionIn` is the description of the parameter which will be printed to the command line tools user if he calls for help and `requiredIn` can be used to force the command line tools user to set a value for this parameter.

The second constructor has the signature (`T& storageIn, const std::string& shortNameIn, const std::string& longNameIn, const std::string& descriptionIn, const T& defaultValueIn`) and is identical to the first constructor except that `requiredIn` is exchanged with `defaultValueIn` to allow the usage of a default value if the command line tools user doesn't set a value for this parameter.

The third constructor has the signature (`T& storageIn, const std::string& shortNameIn, const std::string& longNameIn, const std::string& descriptionIn, const bool requiredIn, const CONTAINER& permittedValuesIn`) and is an expansion of the first constructor extended by the parameter `permittedValuesIn` which is used to provide a container containing all allowed values for the argument.

The fourth constructor has the signature (`T& storageIn, const std::string& shortNameIn, const std::string& longNameIn, const std::string& descriptionIn, const T& defaultValueIn, const CONTAINER& permittedValuesIn`) and is an expansion of the second constructor extended by the parameter `permittedValuesIn` which is used to provide a container containing all allowed values for the argument.

Say you have implemented `mySolver` in a way that the user has the possibility to influence the behaviour of `mySolver`. For instance, the parameter of `mySolver` might have the following options to be set: (i) *x0* (A possible starting point represented by a vector of matching dimension), (ii) *fastMath* (A flag to activate faster mathematic computations), (iii) *eps* (A double value representing the allowed tolerance), (iv) *maxIter* (An integer value representing the maximum number of iterations), (v) *heuristic* (A switch to select the desired heuristic. Allowed values are: `DefaultHeuristic`, `FastHeuristic` or `StrictHeuristic`). Then the caller for `mySolver` with respect to the possible parameter would look as follows:

```

#include <opengm/opengm.hxx> 1
2
// implementation of mySolver 3
#include <opengm/inference/mySolver.hxx> 4
5
// caller base class 6
#include "inference.caller_base.hxx" 7
8
// command line arguments 9
#include "../argument/argument.hxx" 10
11
namespace opengm { 12
namespace interface { 13
14
template <class IO, class GM, class ACC> 15
class mySolverCaller : public InferenceCallerBase<IO, GM, ACC> { 16
protected: 17
    using InferenceCallerBase<IO, GM, ACC>::io_; 18
    using InferenceCallerBase<IO, GM, ACC>::infer; 19

```

```

typedef typename opengm::mySolver<GM, ACC> mySolver;
typename mySolver::Parameter mySolverParameter_;
typedef typename mySolver::TimingVisitorType TimingVisitorType;
virtual void
runImpl(GM& model, StringArgument<>& outputfile, const bool verbose);

std::string selectedHeuristic_;

public:
const static std::string name_;
mySolverCaller(IO& ioIn);
};

template <class IO, class GM, class ACC>
inline mySolverCaller<IO, GM, ACC>::ICMCaller(IO& ioIn)
: InferenceCallerBase<IO, GM, ACC>("MYSOLVER",
" detailed_description_of_mySolver_Parser...", ioIn) {
// add commandline arguments required by mySolver here...
addArgument(VectorArgument<std::vector<size_t> >(mySolverParameter_.startPoint_, "x0",
" startingpoint", " location_of_the_file_containing_the_values_for_the_starting_point", false));
addArgument(BoolArgument(mySolverParameter_.fastMath_, "", " fastMath",
" Allow_functions_to_be_shared"));
addArgument(DoubleArgument<>(mySolverParameter_.eps_, "", " eps",
" allowed_tolerance", mySolverParameter_.eps));
addArgument(Size_TArgument<>(mySolverParameter_.numberOfIterations_, "", " maxIter",
" Maximum_number_of_iterations.", mySolverParameter_.numberOfIterations_));

std::vector<std::string> possibleHeuristics;
possibleHeuristics.push_back(" DefaultHeuristic");
possibleHeuristics.push_back(" FastHeuristic");
possibleHeuristics.push_back(" StrictHeuristic");
addArgument(StringArgument<>(selectedHeuristic_, "", " heuristic",
" Select_desired_heuristic.", possibleHeuristics.front(), possibleHeuristics));
}

template <class IO, class GM, class ACC>
inline void
mySolverCaller<IO, GM, ACC>::runImpl(GM& model,
StringArgument<>& outputfile, const bool verbose) {
std::cout << "running_mySolver_caller" << std::endl;

// add code for setting up mySolver here...
// transform string into corresponding parameter value
if(selectedHeuristic_ == " DefaultHeuristic") {
mySolverParameter_.heuristic_ = mySolver::Parameter::DEFAULT;
} else if(selectedInferenceType_ == " FastHeuristic") {
mySolverParameter_.heuristic_ = mySolver::Parameter::FAST;
} else if(selectedInferenceType_ == " StrictHeuristic") {
mySolverParameter_.heuristic_ = mySolver::Parameter::STRICT;
} else {
throw RuntimeError(" Unknown_inference_type_for_mySolver");
}

// run inference and protocolate results
this-> template infer<ICM, TimingVisitorType,
typename mySolverCaller::Parameter>(model, outputfile, verbose,
mySolverParameter_);
}

template <class IO, class GM, class ACC>
const std::string mySolverCaller<IO, GM, ACC>::name_ = "MYSOLVER";

} // namespace interface
} // namespace opengm

```

## 4.5 Adding Return Types

If protocol mode is active, the results computed by an inference algorithm called by the command line tools are stored using the default timing visitor. This is done in the template method *protocolate()* of the base class *InferenceCallerBase*. If you want to get more information out of your inference call, you have to specialize the visitors or write a new visitor which logs the desired information. Further more the template method *protocolate()* has to be specialized for the new visitor type. The base template method should provide a sufficient example on how the specialized method is supposed to look like.

## Chapter 5

# OpenGM and MATLAB

OpenGM does not come with a full-fledged MATLAB interface. It does, however, contain the C++ code for a MATLAB mex file to construct OpenGM models within MATLAB and save these in the native OpenGM file format. The file format can be read by any OpenGM C++ code, in particular by OpenGM's command line tools. These write results to an HDF5 file that can be read back into MATLAB.

The code below is for a MATLAB mex file that builds second order models. It can be called in MATLAB with matrices `first_order_functions` and `pairs` and a cell-array of matrices `second_order_functions` as `opengmBuild(first_order_functions, pairs, second_order_functions, 'model.h5')`. The reader will find it easy to extend this code to build more complex models.

```
#include "mex.h" 1
#include <stdio.h> 2
#include <stdlib.h> 3
#include <string.h> 4
#include <math.h> 5
#include <vector> 6
7
#include <opengm/operations/adder.hxx> 8
#include <opengm/graphicalmodel/graphicalmodel.hxx> 9
#include <opengm/graphicalmodel/graphicalmodel_hdf5.hxx> 10
#include <opengm/utilities/metaprogramming.hxx> 11
12
typedef double ValueType; 13
typedef size_t IndexType; 14
typedef unsigned char LabelType; 15
typedef opengm::GraphicalModel< 16
    ValueType, 17
    opengm::Adder, 18
    opengm::meta::TypeListGenerator<opengm::ExplicitFunction<ValueType> >::type, 19
    opengm::DiscreteSpace<IndexType, LabelType>, 20
    false 21
    > GraphicalModelType; 22
23
typedef opengm::ExplicitFunction<ValueType> ExplicitFunctionType; 24
typedef GraphicalModelType::FunctionIdentifier FunctionIdentifier; 25
26
void mexFunction(int nlhs, mxArray *plhs[], int nrhs, const mxArray *prhs[]) 27
{ 28
    /* INPUT */ 29
    /* 1.) unaray*/ 30
    /* 2.) pairs*/ 31
    /* 3.) pairfunction*/ 32
    /* 4.) output-filename */ 33
    34
    if (nrhs != 4) { 35
        mexErrMsgTxt(" Incorrect number of inputs."); 36
    }
```

```

}
37
char* filename;
38
filename = (char*)mxCalloc(mxGetN(prhs[3])+1,sizeof(char));
39
40
mxGetString(prhs[3], filename, mxGetN(prhs[3])+1);
41
42
double* unary = (double*) mxGetData(prhs[0]);
43
double* pairwise_node = (double*) mxGetData(prhs[1]);
44
size_t numberOfVariables = mxGetN(prhs[0]);
45
size_t numberOfFactors = mxGetN(prhs[1]);
46
47
/* Output Information*/
48
printf("Number_of_variables:.....%d\n",numberOfVariables);
49
printf("Number_of_pairwise_factors:.....%d\n",numberOfFactors);
50
printf("Output_file:.....%s\n",filename);
51
52
std::vector<LabelType> numStates(numberOfVariables,numberOfVariables);
53
GraphicalModelType gm(
54
    opengm::DiscreteSpace<IndexType, LabelType >(numStates.begin(), numStates.end())
55
);
56
57
//Add Unaries
58
std::cout << "add_Unaries" <<std::endl;
59
for(size_t var=0; var<numberOfVariables; ++var) {
60
    LabelType shape[] = {gm.numberofLabels(var)};
61
    ExplicitFunctionType func(shape,shape+1);
62
    for(size_t i=0; i<shape[0]; ++i) {
63
        func(i) = unary[i+var*numberOfVariables];
64
    }
65
    FunctionIdentifier fid = gm.addSharedFunction(func);
66
    gm.addFactor(fid, &var, &var+1);
67
}
68
69
//ADD Pairwise
70
std::cout << "add_pairwise" <<std::endl;
71
for(size_t f=0;f<numberOfFactors; ++f) {
72
    IndexType vars[2];
73
    LabelType shape[]={numberOfVariables,numberOfVariables};
74
    vars[0] = pairwise_node[0+2*f];
75
    vars[1] = pairwise_node[1+2*f];
76
    ExplicitFunctionType func(shape,shape+2);
77
    mxArray* data = mxGetCell(prhs[2], f);
78
    double* pdata = (double*) mxGetData(data);
79
    for(size_t i1=0; i1<shape[0]; ++i1) {
80
        for(size_t i2=0; i2<shape[1]; ++i2) {
81
            func(i1,i2) = pdata[i1+i2*shape[0]];
82
        }
83
    }
84
    FunctionIdentifier fid = gm.addSharedFunction(func);
85
    gm.addFactor(fid, vars, vars+2);
86
}
87
88
opengm::hdf5::save(gm, filename, "gm");
89
mxFree(filename);
90
}
91

```

## Chapter 6

# Extending OpenGM

### 6.1 Label Spaces

In special cases, neither of the label spaces described in Section 2.3 is maximally efficient. As an example, consider models in which one special variable attains a large number of labels while all other variables are binary. OpenGM accounts for such special cases by allowing the user to write their own label spaces. The custom class template needs to inherit from `SpaceBase`:

```
#include "opengm/opengm.hxx" 1
#include "opengm/graphicalmodel/space/space_base.hxx" 2
template<class I = std::size_t, class L = std::size_t> 3
class CustomSpace : public SpaceBase<CustomSpace<I, L>, I, L> { 4
public: 5
    typedef I IndexType; 6
    typedef L LabelType; 7
    CustomSpace(); 8
    CustomSpace(const IndexType labels, const IndexType variables) 9
        : numberOfLabelsOfFirstVariable_(labels), 10
          numberOfVariables_(variables) 11
    { OPENGM_ASSERT(labels > 0); 12
      OPENGM_ASSERT(variables > 1); } 13
    void assign(const IndexType labels, const IndexType variables) 14
    { OPENGM_ASSERT(labels > 0); 15
      OPENGM_ASSERT(variables > 1); 16
      numberOfLabelsOfFirstVariable_ = labels; 17
      numberOfVariables_ = variables; } 18
    IndexType addVariable(const IndexType labels) 19
    { OPENGM_ASSERT(labels == 2); 20
      ++numberOfVariables_; } 21
    IndexType numberOfVariables() const 22
    { return numberOfVariables_; } 23
    LabelType numberOfLabels(const IndexType variable) const 24
    { return variable == 0 ? numberOfLabelsOfFirstVariable_ : 2; } 25
private: 26
    IndexType numberOfLabelsOfFirstVariable_; 27
    IndexType numberOfVariables_; 28
}; 29
```

### 6.2 Functions

In principle, every graphical model function in OpenGM can be implemented by its value table, using the class `ExplicitFunction`. However, as described in Section 2.4, this is often not the most

efficient representation. Therefore, OpenGM allows the user to write their own class templates for functions.

As an example, consider the modulo function  $f(x) = x \bmod 2$  that yields 0, if  $x$  is even, and 1, if  $x$  is odd. Implementing this function by its value table would surely be a waste of memory if  $x$  can assume many different labels. A class template for the modulo function, as well as any custom function class template, needs to inherit from `FunctionBase`:

```

#include "opengm/opengm.hxx" 1
#include "opengm/functions/function_registration.hxx" 2
#include "opengm/functions/function_properties_base.hxx" 3
template<class T, class I = size_t, class L = size_t> 4
class ModuloFunction 5
: public FunctionBase<ModuloFunction<T, I, L>, T, size_t, size_t> 6
{ 7
public: 8
    typedef T ValueType; 9
    typedef L LabelType; 10
    typedef I IndexType; 11
    ModuloFunction(const IndexType shape) 12
        : shape_(shape); {} 13
    LabelType shape(const size_t j) const 14
    { OPENGM_ASSERT(j == 0); return shape_; } 15
    size_t size() const 16
    { return shape_; } 17
    size_t dimension() const 18
    { return 1; } 19
    template<class ITERATOR> ValueType operator()(ITERATOR it) const 20
    { OPENGM_ASSERT(*it < shape_); return *it % 2; } 21
    bool operator==(const ModuloFunction& other) const 22
    { return other.shape_ == shape_; } 23
    IndexType numberOfParameters() const 24
    { return 1; } 25
    IndexType parameter(const size_t j) const 26
    { OPENGM_ASSERT(j == 0); return shape_; } 27
    IndexType& parameter(const size_t j); 28
    { OPENGM_ASSERT(j == 0); return shape_; } 29
private: 30
    IndexType shape_; 31
friend class FunctionSerialization<ModuloFunction<T, I, L> >; 32
}; 33

```

Every custom function needs to be registered with a unique ID between 0 and 16000 (IDs above 16000 are reserved for the OpenGM standard). Registration works as follows:

```

template <class T, class I, class L> 1
struct FunctionRegistration<ModuloFunction<T, I, L> > { 2
    enum ID { Id = 0; }; 3
}; 4

```

One last step is necessary to add the custom class of functions to the OpenGM file format. This works by specializing the class template `FunctionSerialization`:

```

template<class T, class I, class L> 1
class FunctionSerialization<ModuloFunction<T, I, L> > { 2
public: 3
    typedef typename ModuloFunction<T, I, L>::ValueType ValueType; 4
    static size_t indexSequenceSize(const ModuloFunction<T, I, L>&) 5
    { return 1; } 6
    static size_t valueSequenceSize(const ModuloFunction<T, I, L>&) 7

```



```

    { return 0; }
template<class INDEX_ITERATOR, class VALUE_ITERATOR>
    static void serialize(const ModuloFunction<T, I, L>& f,
        INDEX_ITERATOR iit, VALUE_ITERATOR vit)
    { *iit = f.size(); }
template<class INDEX_ITERATOR, class VALUE_ITERATOR>
    static void deserialize(INDEX_ITERATOR iit, VALUE_ITERATOR vit,
        ModuloFunction<T, I, L>& f)
    { f.shape_ = *iit; }
};

```

## 6.3 Operations

Commutative and associative operations are classes in OpenGM, equipped with member functions for the operation itself and its inverse, the neutral and inverse neutral element. Member functions for an (inverse) hyper-operation and an order associated with the operation are optional. The complete interface (for brevity without `static` and `const` qualifiers) is shown in Tab. 6.1. Operations are used in algorithms like in the example below:

<code>T neutral&lt;T&gt;()</code>	returns the neutral element
<code>void neutral(T&amp; out)</code>	returns the neutral element by reference
<code>T inneutral&lt;T&gt;()</code>	returns the inverse neutral element
<code>void inneutral(T&amp; out)</code>	returns the inverse neutral element by reference
<code>void op(T1&amp; in1, T2&amp; out)</code>	in-place operation ( <code>out = op(out,in1)</code> )
<code>void op(T1&amp; in1, T2&amp; in2, T3&amp; out)</code>	operation ( <code>out =op(in1,in2)</code> )
<code>void iop(T1&amp; in1, T2&amp; out)</code>	in-place inverse operation ( <code>out = iop(out,in1)</code> )
<code>void iop(T1&amp; in1, T2&amp; in2, T3&amp; out)</code>	inverse operation ( <code>out = iop(in1,in2)</code> )
<code>void hop( T1&amp; in, T2&amp; out)</code>	in-place hyper-operation ( <code>out = hop(out,in1)</code> )
<code>void hop( T1&amp; in1,T2&amp; in2, T3&amp; out)</code>	hyper-operation ( <code>out =hop(in1,in2)</code> )
<code>void ihop(T1&amp; in, T2&amp; out)</code>	in-place inverse hyper-operation ( <code>out = ihop(out,in1)</code> )
<code>void ihop(T1&amp; in1,T2&amp; in2, T3&amp; out)</code>	inverse hyper-operation ( <code>out = ihop(in1,in2)</code> )
<code>bool hasbop()</code>	returns true if operation provides a compare operation
<code>bool bop(T&amp; in1, T&amp; in2)</code>	boolean compare operation
<code>bool ibop(T&amp; in1, T&amp; in2)</code>	inverse boolean compare operation

Table 6.1: Interface of operations in OpenGM

```

#include <opengm/datastructures/marray/marray.hxx>
#include <opengm/operations/adder.hxx>
typedef double T;
typedef opengm::Adder OP;
std::vector<size_t> shape(2,4);
marray::Marray<T> A(shape.begin(),shape.end(),2);
marray::Marray<T> B(shape.begin(),shape.end(),6);
marray::Marray<T> C(shape.begin(),shape.end(),1);
OP::op(A,B,C); //C(..)=8
OP::iop(A,C); //C(..)=10
double d = OP::neutral<double>(); //d=0
OP::ineutral(d); //d=oo

```

## 6.4 Algorithms

All algorithms implemented in OpenGM access graphical models via the general interface described in Section 2.8. This interface is designed w.r.t. performance because its functions are called often,

in inner loops of algorithms<sup>1</sup>.

Algorithms themselves share a parsimonious interface called `Inference`<sup>2</sup>. Deriving a custom algorithm from `Inference` has advantages, e.g. when it comes to using algorithms exchangeably or adding custom algorithms to the OpenGM command line tools (Section 4).

By far the best example of how to implement a custom algorithm is the simple brute force search, `include/opengm/inference/bruteforce.hxx`.

---

<sup>1</sup>Inheritance with virtual functions is avoided altogether within a graphical model.

<sup>2</sup>Inheritance with virtual functions in the interface class template `Inference` is used such that algorithms are exchangeable. The effect on the runtime is negligible because the interface of *algorithms* will almost never become a runtime bottleneck.

# Chapter 7

## Examples

### 7.1 Markov Chain

```
#include <opengm/graphicalmodel/graphicalmodel.hxx> 1
#include <opengm/graphicalmodel/space/simplediscretespace.hxx> 2
#include <opengm/functions/potts.hxx> 3
#include <opengm/operations/adder.hxx> 4
#include <opengm/inference/messagepassing.hxx> 5
6
using namespace std; // 'using' is used only in example code 7
using namespace opengm; 8
9
int main() { 10
    // construct a label space with numberOfVariables many variables,
    // each having numberOfLabels many labels 11
    const size_t numberOfVariables = 40; 12
    const size_t numberOfLabels = 5; 13
    typedef SimpleDiscreteSpace<size_t, size_t> Space; 14
    Space space(numberOfVariables, numberOfLabels); 15
16
    // construct a graphical model with 17
    // - addition as the operation (template parameter Adder) 18
    // - support for Potts functions (template parameter PottsFunction<double>) 19
    typedef OPENGM_TYPERELIST_2(ExplicitFunction<double>, PottsFunction<double>) FunctionTypelist; 20
    typedef GraphicalModel<double, Adder, FunctionTypelist, Space> Model; 21
    Model gm(space); 22
23
    // for each variable, add one 1st order functions and one 1st order factor 24
    for(size_t v = 0; v < numberOfVariables; ++v) { 25
        const size_t shape[] = {numberOfLabels}; 26
        ExplicitFunction<double> f(shape, shape + 1); 27
        for(size_t s = 0; s < numberOfLabels; ++s) { 28
            f(s) = static_cast<double>(rand()) / RAND_MAX; 29
        } 30
        Model::FunctionIdentifier fid = gm.addFunction(f); 31
    } 32
    size_t variableIndices[] = {v}; 33
    gm.addFactor(fid, variableIndices, variableIndices + 1); 34
} 35
36
// add one (!) 2nd order Potts function 37
PottsFunction<double> f(numberOfLabels, numberOfLabels, 0.0, 0.3); 38
Model::FunctionIdentifier fid = gm.addFunction(f); 39
40
// for each pair of consecutive variables, 41
// add one factor that refers to the Potts function 42
for(size_t v = 0; v < numberOfVariables - 1; ++v) { 43
    size_t variableIndices[] = {v, v + 1}; 44
} 45
```

```

    gm.addFactor(fid, variableIndices, variableIndices + 2);           46
}                                                                      47
// set up the optimizer (loopy belief propagation)                    48
// set up the optimizer (loopy belief propagation)                    49
typedef BeliefPropagationUpdateRules<Model, Minimizer> UpdateRules;    50
typedef MessagePassing<Model, Minimizer, UpdateRules, MaxDistance> BeliefPropagation; 51
const size_t maxNumberOfIterations = numberOfVariables * 2;          52
const double convergenceBound = 1e-7;                                53
const double damping = 0.0;                                          54
BeliefPropagation::Parameter parameter(maxNumberOfIterations, convergenceBound, damping); 55
BeliefPropagation bp(gm, parameter);                                  56
// optimize (approximately)                                          57
BeliefPropagation::VerboseVisitorType visitor;                       58
bp.infer(visitor);                                                  59
// obtain the (approximate) argmin                                    60
vector<size_t> labeling(numberOfVariables);                          61
bp.arg(labeling);                                                  62
}                                                                      63
}                                                                      64
}                                                                      65

```

## 7.2 Bipartite Matching

```

#include <iostream>                                                    1
#include <iomanip>                                                      2
#include <vector>                                                       3
#include <string>                                                       4
                                                                    5
#include <opengm/opengm.hxx>                                           6
#include <opengm/datastructures/marray/marray.hxx>                     7
#include <opengm/functions/potts.hxx>                                  8
#include <opengm/graphicalmodel/graphicalmodel.hxx>                   9
#include <opengm/operations/adder.hxx>                                 10
#include <opengm/inference/astar.hxx>                                  11
                                                                    12
using namespace std; // 'using' is used only in example code         13
using namespace opengm;                                              14
                                                                    15
template<class T>                                                     16
void createAndPrintData(size_t nrOfVariables, marray::Marray<T> & data) 17
{
    size_t shape[]={nrOfVariables,nrOfVariables};                    18
    data.resize(shape,shape+2);                                       19
    cout<<" pairwise_costs:\n";                                       20
    srand(0);                                                         21
    for(size_t v=0;v<data.shape(0);++v) {                             22
        for(size_t s=0;s<data.shape(0);++s) {                         23
            data(v,s)=static_cast<float>(rand()%100) *0.01;          24
            cout<<left<<setw(6)<<setprecision(2)<<data(v,s);          25
        }                                                             26
        cout<<"\n";                                                  27
    }
}
}
void printSolution(const vector<size_t> & solution)                    28
{
    set<size_t> unique;                                               29
    cout<<"\nSolution_States_:\n";                                     30
    for(size_t v=0;v<solution.size();++v) {                           31
        cout<<left<<setw(2)<<v<<"_>_><<solution[v]<<"\n";          32
    }
}
}
int main(int argc, char** argv)                                       33
{
    // model parameters                                              34
    const size_t nrOfVariables=5;                                     35
}
}

```

```

    const size_t nrOfStates=nrOfVariables;
float high=20;
cout<<"\n_Matching_with_one_to_one_correspondences:\n"
<<nrOfVariables <<"variables_with_"<<nrOfStates<<"_states_\n\n";
// create data terms / pairwise costs
marray::Marray<float> data;
createAndPrintData(nrOfVariables,data);
// Build the Model
// type of the graphical model:
typedef GraphicalModel
<
    float,
    Adder,
    OPENGM_TYPELIST_2(PottsFunction<float>,opengm::ExplicitFunction<float>)
>
Model;
typedef ExplicitFunction<float> ExplicitFunctionType ;
typedef Model::FunctionIdentifier FunctionIdentifier;

// construct a graphical model with
// - nrOfVariables variables,
// each having nrOfStates labels
// - addition as the operation (template parameter Adder)
// - support for Potts functions (template parameter PottsFunction<double>))

vector<size_t> variableStates(nrOfVariables,nrOfStates);
Model gm(DiscreteSpace<size_t,size_t> (variableStates.begin(),variableStates.end()));
//unary potentials
{
    //shape of the function f
    const size_t shape[]={nrOfStates};
    //construct the function
    ExplicitFunctionType f(shape,shape+1);
    //add factors and functions
    for(size_t v=0;v<nrOfVariables;++v) {
        for(size_t s=0;s<nrOfStates;++s) {
            f(s)=1.0f-data(v,s);
        }
        FunctionIdentifier id=gm.addFunction(f);
        size_t vi[]={v};
        gm.addFactor(id,vi,vi+1);
    }
}
//pair potentials (Potts)
{
    // add one (!) 2nd order Potts function
    PottsFunction<float> f(nrOfStates,nrOfStates,high,0);
    FunctionIdentifier id=gm.addFunction(f);
    //add pair potentials for all variables
    for(size_t v1=0;v1<nrOfVariables;++v1)
    for(size_t v2=v1+1;v2<nrOfVariables;++v2) {
        //create the variables indices for the factor
        size_t vi[]={ v1,v2 };
        // sort the variables indices is not necessary
        // since v1 is v1 < v2 true by construction
        // but remember SORTED VARIABLE SEQUENCE IS PRECODITION
        // add factor
        gm.addFactor(id,vi,vi+2);
    }
}
// set up the optimizer (A*Star)
typedef AStar<Model, Minimizer > AstarType;
AstarType astar(gm);

// obtain the argmin
AstarType::VerboseVisitorType verboseVisitor;
cout<<"\nStart_Astar:\n";

```

```

    astar.infer(verboseVisitor);
    // output the argmin
    vector<size_t> solution;
    astar.arg(solution);
    printSolution(solution);
}

```

### 7.3 Interpixel-Boundary Segmentation

```

#include <iostream>
#include <iomanip>
#include <vector>
#include <string>
#include <algorithm>
#include <cmath>

#include <opengm/opengm.hxx>
#include <opengm/datastructures/marray/marray.hxx>
#include <opengm/graphicalmodel/space/simplifiediscretespace.hxx>
#include <opengm/graphicalmodel/graphicalmodel.hxx>
#include <opengm/operations/adder.hxx>
#include <opengm/inference/lazyflipper.hxx>

using namespace std; // 'using' is used only in example code

// this class is used to map a node (x, y) in the topological
// grid to a unique variable index
class TopologicalCoordinateToIndex {
public:
    TopologicalCoordinateToIndex(
        const size_t geometricGridSizeX,
        const size_t geometricGridSizeY
    )
    : gridSizeX_(geometricGridSizeX),
      gridSizeY_(geometricGridSizeY)
    {}

    const size_t operator()(
        const size_t tx,
        const size_t ty
    ) const {
        return tx / 2 + (ty / 2)*(gridSizeX_) + ((ty + ty % 2) / 2)*(gridSizeX_ - 1);
    }

    size_t gridSizeX_;
    size_t gridSizeY_;
};

template<class T>
void randomData(
    const size_t gridSizeX,
    const size_t gridSizeY,
    marray::Marray<T>& data
) {
    srand(gridSizeX * gridSizeY);
    const size_t shape[] = {gridSizeX, gridSizeY};
    data.assign();
    data.resize(shape, shape + 2);
    for (size_t y = 0; y < gridSizeY; ++y) {
        for (size_t x = 0; x < gridSizeX; ++x) {
            data(x, y) = static_cast<float>(rand() % 10) * 0.1;
        }
    }
}

```

```

56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122

template<class T>
void printData(
    const marray::Marray<T> & data
) {
    cout << "energy_for_boundary_to_be_active:" << endl;
    for (size_t y = 0; y < data.shape(1)*2 - 1; ++y) {
        for (size_t x = 0; x < data.shape(0)*2 - 1; ++x) {
            if (x % 2 == 0 && y % 2 == 0) {
                cout << left << setw(3) << setprecision(1) << data(x / 2, y / 2);
            } else if (x % 2 == 0 && y % 2 == 1) {
                cout << left << setw(3) << setprecision(1) << "___";
            } else if (x % 2 == 1 && y % 2 == 0) {
                cout << left << setw(3) << setprecision(1) << "._.";
            } else if (x % 2 == 1 && y % 2 == 1) {
                cout << left << setw(3) << setprecision(1) << "._.";
            }
        }
        cout << "\n";
    }
}

// output the (approximate) argmin
template<class T>
void printSolution(
    const marray::Marray<T>& data,
    const vector<size_t>& solution
) {
    TopologicalCoordinateToIndex cTHelper(data.shape(0), data.shape(1));
    cout << endl << "solution_states:" << endl;
    cout << "solution:" << endl;
    for (size_t x = 0; x < data.shape(0)*2 - 1; ++x) {
        cout << left << setw(3) << setprecision(1) << "___";
    }
    cout << endl;
    for (size_t y = 0; y < data.shape(1)*2 - 1; ++y) {
        cout << "|";
        for (size_t x = 0; x < data.shape(0)*2 - 1; ++x) {
            if (x % 2 == 0 && y % 2 == 0) {
                data(x / 2, y / 2) = static_cast<float>(rand() % 10) * 0.1;
                cout << left << setw(3) << setprecision(1) << "._.";
            } else if (x % 2 == 0 && y % 2 == 1) {
                if (solution[cTHelper(x, y)]) {
                    cout << left << setw(3) << setprecision(1) << "___";
                }
                else {
                    cout << left << setw(3) << setprecision(1) << "._.";
                }
            } else if (x % 2 == 1 && y % 2 == 0) {
                if (solution[cTHelper(x, y)])
                    cout << left << setw(3) << setprecision(1) << "._.";
                else
                    cout << left << setw(3) << setprecision(1) << "___";
            } else if (x % 2 == 1 && y % 2 == 1) {
                cout << left << setw(3) << setprecision(1) << "._.";
            }
        }
        cout << "|" << endl;
    }
    for (size_t x = 0; x < data.shape(0)*2 - 1; ++x) {
        cout << left << setw(3) << setprecision(1) << "___";
    }
    cout << endl;
}

// user defined Function Type

```

```

template<class T> 123
struct ClosednessFunctor { 124
public: 125
    typedef T value_type; 126
    127
    template<class Iterator> 128
    inline const T operator()(Iterator begin)const { 129
        size_t sum = begin[0]; 130
        sum += begin[1]; 131
        sum += begin[2]; 132
        sum += begin[3]; 133
        if (sum != 2 && sum != 0) { 134
            return high; 135
        } 136
        return 0; 137
    }; 138
    size_t dimension()const { 139
        return 4; 140
    }; 141
    size_t shape(const size_t i)const { 142
        return 2; 143
    }; 144
    size_t size()const { 145
        return 16; 146
    }; 147
    T high; 148
}; 149
150
int main(int argc, char** argv) { 151
    // model parameters 152
    const size_t gridSizeX = 5, gridSizeY = 5; //size of grid 153
    const float beta = 0.9; // bias to choose between under- and over-segmentation 154
    const float high = 10; // closedness-enforcing soft-constraint 155
    156
    // size of the topological grid 157
    const size_t tGridSizeX = 2 * gridSizeX - 1, tGridSizeY = 2 * gridSizeY - 1; 158
    const size_t nrOfVariables = gridSizeX * (gridSizeX - 1) + gridSizeX * (gridSizeY - 1); 159
    const size_t dimT[] = {tGridSizeX, tGridSizeY}; 160
    TopologicalCoordinateToIndex cTHelper(gridSizeX, gridSizeY); 161
    marray::Marray<float> data; 162
    randomData(gridSizeX, gridSizeY, data); 163
    164
    cout << "interpixel_boundary_segmentation_with_closedness:" << endl; 165
    printData(data); 166
    167
    // construct a graphical model with 168
    // - addition as the operation (template parameter Adder) 169
    // - the user defined function type ClosednessFunctor<float> 170
    // - gridSizeY * (gridSizeX - 1) + gridSizeX * (gridSizeY - 1) variables, 171
    // each having 2 many labels. 172
    typedef opengm::meta::TypeListGenerator< 173
        opengm::ExplicitFunction<float>, 174
        ClosednessFunctor<float> 175
    >::type FunctionTypeList; 176
    typedef opengm::GraphicalModel<float, opengm::Adder, FunctionTypeList, 177
        opengm::SimpleDiscreteSpace<> > Model; 178
    typedef Model::FunctionIdentifier FunctionIdentifier; 179
    Model gm(opengm::SimpleDiscreteSpace<>(nrOfVariables, 2)); 180
    181
    // for each boundary in the grid, i.e. for each variable 182
    // of the model, add one 1st order functions 183
    // and one 1st order factor 184
    { 185
        const size_t shape[] = {2}; 186
    } 187
}; 188
189

```



```

opengm::ExplicitFunction<float> f(shape, shape + 1);
for (size_t yT = 0; yT < dimT[1]; ++yT) {
  for (size_t xT = 0; xT < dimT[0]; ++xT) {
    if ((xT % 2 + yT % 2) == 1) {
      float gradient = fabs(data(xT / 2, yT / 2) - data(xT / 2 + xT % 2, yT / 2 + yT % 2));
      f(0) = beta * gradient; // value for inactive boundary
      f(1) = (1.0 - beta) * (1.0 - gradient); // value for active boundary;
      FunctionIdentifier id = gm.addFunction(f);
      size_t vi[] = {cTHelper(xT, yT)};
      gm.addFactor(id, vi, vi + 1);
    }
  }
}

// for each junction of four inter-pixel edges on the grid,
// one factor is added that connects the corresponding variable
// indices and refers to the ClosednessFunctor function
{
  // add one (!) 4th order ClosednessFunctor function
  ClosednessFunctor<float> f;
  f.high = high;
  FunctionIdentifier id = gm.addFunction(f);
  // add factors
  for (size_t y = 0; y < dimT[1]; ++y) {
    for (size_t x = 0; x < dimT[0]; ++x) {
      if (x % 2 + y % 2 == 2) {
        size_t vi[] = {
          cTHelper(x + 1, y),
          cTHelper(x - 1, y),
          cTHelper(x, y + 1),
          cTHelper(x, y - 1)
        };
        sort(vi, vi + 4);
        gm.addFactor(id, vi, vi + 4);
      }
    }
  }
}

// set up the optimizer (lazy flipper)
typedef opengm::LazyFlipper<Model, opengm::Minimizer> LazyFlipperType;
LazyFlipperType::VerboseVisitorType verboseVisitor;
size_t maxSubgraphSize = 5;
LazyFlipperType lazyflipper(gm, maxSubgraphSize);
cout << "start_inference:" << endl;

// obtain the (approximate) argmin
lazyflipper.infer(verboseVisitor);

// output the (approximate) argmin
vector<size_t> solution;
lazyflipper.arg(solution);
printSolution(data, solution);
}

```



## Chapter 8

# License

Copyright (C) 2012 Bjoern Andres, Thorsten Beier and Joerg H. Kappes.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.



# Bibliography

- [1] B. Andres, J. H. Kappes, U. Koethe, and F. A. Hamprecht. The Lazy Flipper: MAP inference in higher-order graphical models by depth-limited exhaustive search. *ArXiv e-prints*, 2010.
- [2] B. Andres, T. Koethe, U. Kroeger, and F. A. Hamprecht. Runtime-flexible multi-dimensional arrays and views for C++98 and C++0x. *ArXiv e-prints*, Aug. 2010.
- [3] A. Barbu and S.-C. Zhu. Generalizing swendsen-wang to sampling arbitrary posterior probabilities. *TPAMI*, 27(8):1239–1253, 2005.
- [4] M. Bergtholdt, J. H. Kappes, S. Schmidt, and C. Schnörr. A study of parts-based object class detection using complete graphs. *International Journal of Computer Vision*, 87(1-2):93–117, 2010.
- [5] J. Besag. On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society B*, 48:259–302, 1986.
- [6] E. Boros, P. L. Hammer, and X. Sun. Network flows and minimization of quadratic pseudo-boolean functions. Technical report, Technical Report RRR 17-1991, RUTCOR Research Report, 1991.
- [7] E. Boros, P. L. Hammer, and G. Tavares. Preprocessing of unconstrained quadratic binary optimization. Technical report, Technical Report RRR 10-2006, RUTCOR Research Report, 2006.
- [8] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *TPAMI*, 23(11):1222–1239, 2001.
- [9] S. Geman and D. Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *TPAMI*, 6:721–741, 1984.
- [10] A. V. Goldberg, S. Hed, H. Kaplan, R. E. Tarjan, and R. F. Werneck. Maximum flows by incremental breadth-first search. In *Proceedings of the 19th European conference on Algorithms*, ESA’11, pages 457–468, Berlin, Heidelberg, 2011. Springer-Verlag.
- [11] P. Hammer, P. Hansen, and B. Simeone. Roof duality, complementation and persistency in quadratic 0-1 optimization. *Math. Programming*, 28:121–155, 1984.
- [12] HDF5 Group. HDF5.  
[urlhttp://www.hdfgroup.org/HDF5/](http://www.hdfgroup.org/HDF5/).
- [13] S. Hed. IBFS.  
[urlhttp://www.cs.tau.ac.il/~sagihed/ibfs/](http://www.cs.tau.ac.il/~sagihed/ibfs/).
- [14] C. Helmberg. The ConicBundle Library for Convex Optimization v0.3.11.  
[urlhttp://www-user.tu-chemnitz.de/~helmberg/ConicBundle/](http://www-user.tu-chemnitz.de/~helmberg/ConicBundle/), 2012.
- [15] IBM. IBM ILOG CPLEX Optimizer.  
[urlhttp://www-01.ibm.com/software/integration/optimization/cplex-optimizer/](http://www-01.ibm.com/software/integration/optimization/cplex-optimizer/), 2012.

- [16] Jeremy G. Siek et.al.. Boost C++ Libraries v1.47.0.  
url<http://www.boost.org/>, 2011.
- [17] K. Jung, P. Kohli, and D. Shah. Local rules for global map: When do they work ? In *NIPS*, 2009.
- [18] J. H. Kappes. *Inference on Highly-Connected Discrete Graphical Models with Applications to Visual Object Recognition*. doctoral thesis, Ruprecht-Karls-Universität Heidelberg, Faculty of Mathematics and Computer Sciences, Heidelberg, Germany, 2011.
- [19] J. H. Kappes, B. Savchynskyy, and C. Schnörr. A bundle approach to efficient MAP-inference by Lagrangian relaxation. In *CVPR 2012*, 2012. in press.
- [20] P. Kohli, M. P. Kumar, and P. H. S. Torr. P3 & beyond: Solving energies with higher order cliques. In *CVPR*, 2007.
- [21] V. Kolmogorov. Maxflow v3.02.  
url<http://pub.ist.ac.at/vnk/software.html>.
- [22] V. Kolmogorov. QPBO v1.3.  
url<http://pub.ist.ac.at/vnk/software.html>.
- [23] V. Kolmogorov. TRW-S v1.3.  
url<http://research.microsoft.com/en-us/downloads/dad6c31e-2c04-471f-b724-ded18bf70fe3/>.
- [24] V. Kolmogorov. Convergent tree-reweighted message passing for energy minimization. *TPAMI*, 28(10):1568–1583, 2006.
- [25] N. Komodakis, N. Paragios, and G. Tziritas. MRF energy minimization and beyond via dual decomposition. *TPAMI*, 33(3):531–552, 2011.
- [26] F. R. Kschischang, B. J. Frey, and H. Loeliger. Factor graphs and the sum-product algorithm. *Transactions on Information Theory*, 47:498–519, 2001.
- [27] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *Journal of Chemical Physics*, 21(6):1087–1092, 1953.
- [28] J. M. Mooij. libDAI v0.3.0.  
url<http://cs.ru.nl/jorism/libDAI/>.
- [29] J. M. Mooij. libDAI: A free and open source C++ library for discrete approximate inference in graphical models. *JMLR*, 11:2169–2173, Aug. 2010.
- [30] K. Murphy, Y. Weiss, and M. Jordan. Loopy belief propagation for approximate inference: An empirical study. In *UAI*, 1999.
- [31] J. Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann, San Francisco, CA, USA, 1988.
- [32] C. Rother, V. Kolmogorov, V. S. Lempitsky, and M. Szummer. Optimizing binary mrfs via extended roof duality. In *CVPR*, 2007.
- [33] R. Szeliski, R. Zabih, D. Scharstein, O. Veksler, V. Kolmogorov, A. Agarwala, M. Tappen, and C. Rother. MRF-LIB 2.1.  
url<http://vision.middlebury.edu/MRF/code/>.
- [34] R. Szeliski, R. Zabih, D. Scharstein, O. Veksler, V. Kolmogorov, A. Agarwala, M. Tappen, and C. Rother. A comparative study of energy minimization methods for markov random fields with smoothness-based priors. *TPAMI*, 30(6):1068–1080, 2008.

- [35] M. J. Wainwright and M. I. Jordan. *Graphical Models, Exponential Families, and Variational Inference*. Now Publishers Inc., Hanover, MA, USA, 2008.