



The Future of Logic: Foundation-Independence

Florian Rabe

Abstract. Throughout the twentieth century, the automation of formal logics in computers has created unprecedented potential for practical applications of logic—most prominently the mechanical verification of mathematics and software. But the high cost of these applications makes them infeasible but for a few flagship projects, and even those are negligible compared to the ever-rising needs for verification. One of the biggest challenges in the future of logic will be to enable applications at much larger scales and simultaneously at much lower costs. This will require a far more efficient allocation of resources. Wherever possible, theoretical and practical results must be formulated generically so that they can be instantiated to arbitrary logics; this will allow reusing results in the face of today’s multitude of application-oriented and therefore diverging logical systems. Moreover, the software engineering problems concerning automation support must be decoupled from the theoretical problems of designing logics and calculi; this will allow researchers outside or at the fringe of logic to contribute scalable logic-independent tools. Anticipating these needs, the author has developed the MMT framework. It offers a modern approach towards defining, analyzing, implementing, and applying logics that focuses on modular design and logic-independent results. This paper summarizes the ideas behind and the results about MMT. It focuses on showing how MMT provides a theoretical and practical framework for the future of logic.

Mathematics Subject Classification. 03B70, 03B22, 03B15.

Keywords. MMT, Logical declarations, Theory, Theory morphism, Syntax, Model theory, Proof theory, Parsing, Presentation, Type reconstruction, Knowledge management, User interface, Search.

The author was supported by DFG Grant RA-18723-1 OAF.

This paper won the prize of the UNILOG’2015 contest, *The Future of Logic*: <http://www.uni-log.org/future-of-logic>.

1. Introduction and Related Work

Motivation. While logic has historically focused on the theoretical study of a few individual logics—mostly first-order logic and some others such as higher-order or modal logic—recent decades have seen increasing specialization into a plethora of different logics. Moreover, during that time advances in technology—e.g., the internet and theorem provers—have dramatically changed the scale of practical applications of logic. For the future, today’s logicians envision the routine use of logic for the verification of mathematical theorems and safety-critical software. For example, the verification of on-chip algorithms can prevent billion-dollar mishaps such as the Pentium division bug.

However, these approaches pay off almost exclusively at large scales due to the high level of theoretical understanding and practical investment that they require from both developers and users. For example, flagship projects such as the verifications of the Kepler conjecture [16] or the L4 microkernel [28] required double-digit person years of investment.

These scales bring a new set of critical challenges for logics. For instance, they require building large libraries of logical theorems collaboratively, using multiple logics in the same project while reusing theorems across logics, and the interoperability of logic-based tools. Many of these challenges were not anticipated in the designs of current logics, tools, and libraries:

- Modern logic tools such as ACL2 [27], Coq [9], HOL [15, 18, 19, 32], Matita [2], Mizar [44], Nuprl [8], or PVS [33] are built on **mutually incompatible logical foundations**. These include first-order logic, higher-order logic, set theory, and numerous variants of type theory.
- The respective communities are mostly disjoint and have built **overlapping but mutually incompatible libraries**.
- All of these tools **lack urgently-needed features** because no community can afford developing (and maintaining) all desirable features. These include improvements of core components like the theorem prover as well as of peripheral components like user interface or library management.
- Novel logics usually have to start from scratch because it is almost impossible to reuse existing tools and libraries. That massively **slows down evolution** because it can take years to evaluate a new idea.
- All but a few logics are **never used beyond toy examples** because it takes each system dozens of person-years to reach production-ready maturity.

Existing Approaches. These problems have been known for several decades (see, e.g., [1]) and have motivated **three major independent developments in different, disjoint communities**:

(1) **Logical frameworks** [35] are meta-logics in which logics can be defined. Their key benefit is that results about the meta-logic can be inherited by the logics defined in them. Examples include type inference (the Twelf tool [36] for LF [17]), rewriting (the Dedukti tool [4] for LF modulo [10]), and theorem proving (the Isabelle tool [34] for higher-order logic [6]).

But logical frameworks **do not go far enough**: Like logics, the various meta-logics use mutually incompatible foundations and tools. Moreover, they

Mathematics	Logic	Logical framework	Foundation-independence
			MMT
		Meta-logic	Meta-logic
	Logic	Logic	Logic
Domain knowledge	Domain knowledge	Domain knowledge	Domain knowledge

are not expressive enough for defining modern practical logics such as the ones cited above. This requires the design of more and more complex meta-logics, partially defeating the purpose for which they were introduced.

(2) Categorical frameworks like **institutions** [14] give an abstract, uniform definition of what a logic is. Their key benefit is that they capture common concepts of logics (such as theories and models) concisely and leverage category theory to elegantly formalize reuse across theories and logics. Many theoretical results that were originally logic-specific have been generalized to the institution-independent level [12], and institution-independent practical tool support has been developed [31].

But institutions **go too far**: They abstract from the syntactic structure of sentences, theories, and proofs that is essential to logic. Thus, they cannot provide comprehensive logic-independent tool support and remain dependent on specific tools for each logic.

(3) **Markup languages for formal knowledge** such as MathML [3], OpenMath [5], and OMDoc [29] allow the representation of syntax trees for logical objects (e.g., theories or formulas). Their key benefit is that they provide a standardized machine-readable interchange format that can be used for any logic. A major achievement is the inclusion of MathML into the definition of HTML5.

But markup languages succeeded only because they **ignore the logical semantics**. They allow the representation of proofs and models, but they do not make any attempt to define which proofs or models are correct.

Foundation-Independence. We introduce the novel paradigm of *foundation-independence*—the systematic abstraction from logic-specific conceptualizations, theorems, and algorithms in order to obtain results that apply to any logic. This yields the flexibility to design new logics on demand and instantiate existing results to apply them right away.

Contrary to logical frameworks, we do not fix any foundation, not even a meta-logic. Contrary to institutions, we can obtain reusable results that make use of the syntactic structure. And contrary to markup languages, the logical semantics is formalized and respected by the tools.

Systematically aiming for foundation-independence has led the author to develop MMT, which consists of a language [40, 43] and a tool [38]. They provide the theoretical and practical environment to define concepts, reason meta-logically, and implement tool support in a foundation-independent way.

MMT stands for meta-meta-theory/tool. This name is motivated by seeing MMT as part of a **series of abstractions steps** as pictured above. In conventional mathematics, domain knowledge was expressed directly in mathematical notation. Starting with Frege’s work, logic provided a formal syntax and semantics for this notation. Starting in the 1970s, logical frameworks provided meta-logics to formally define this syntax and semantics. Now MMT formalizes the foundation-independent level.

One might expect that this meta-meta-level, at which MMT works, is too abstract to develop deep results. **But not only is it possible to generalize many existing results to the foundation-independent level, MMT-based solutions can even be simpler and stronger than foundation-specific ones.** Moreover, MMT is very well-suited for modularity and system integration and thus better prepared for the large scale challenges of the future than any foundation-specific system. In particular, it systematically separates concerns between logic designers, tool developers, and application developers.

We argue this point by surveying the author’s work of the past 10 years. This includes numerous foundation-independent results, which we summarize in a coherent setting:

- Section 2: the MMT language, which provides the foundation-independent representation of any logical object,
- Section 3: the foundation-independent definition of logic-related concepts based on the MMT language,
- Section 4: the foundation-independent logical algorithms implemented in the MMT tool,
- Section 5: the knowledge management support developed foundation-independently in the MMT tool.

In all cases, we pay special attention to the foundation-independent nature of the results and discuss the differences from and benefits over foundation-specific approaches.

2. Foundation-Independent Representation

Key Concepts. The MMT language uses a small set of carefully chosen orthogonal primitive concepts: *theories*, *symbols*, and *objects*, which are related by *typing* and *equality* and acted on by theory *morphisms*. The main insight behind foundation-independence is that these concepts are at the same time

- *universal* in the sense that they occur in virtually all formal systems in essentially the same way,
- *complete* in the sense that they allow representing virtually all logics and related languages.

Our experiments show that MMT **theories** subsume any kind of formal system such as logical frameworks, mathematical foundations (e.g., ZF set theory), type theories, logics, domain theories, ontology languages, ontologies, specification languages, specifications, etc.

Every theory consists of a list of **symbol declarations** $c[: t][= d][\#N]$ where c is the symbol identifier, the objects t and d are its type and definiens, and N is its notation. Symbol declarations subsume any kind of basic declaration common in formal systems such as type/constant/function/predicate symbols, binders, type operators, concepts, relations, axioms, theorems, inference rules, derived rules, etc. In particular, theorems are just a special case of typed symbols: They can be represented via the propositions-as-types correspondence [11, 21] as declarations $c : F = p$, which establish theorem F via proof p .

MMT **objects** subsume any kind of complex expressions common in formal systems such as terms, values, types, literals, individuals, universes, formulas, proofs, derivations, etc. The key property of theories is that they define the scope of identifiers: An object o over a theory T may use only the symbols visible to T .

Objects o and o' over a theory T are subject to the **typing** and **equality** judgments $\vdash_T o : o'$ and $\vdash_T o \equiv o'$. These judgments subsume any kind of classification used to define the semantics of formal systems: typing subsumes, e.g., well-formedness, typing, kinding, sorting, instantiation, satisfaction, and proving; equality subsumes, e.g., axiomatic equality, platonic equality, rewriting, and computation.

The central step to represent a specific foundation in MMT is (i) to give a theory F that declares one symbol for each primitive operator of the foundation, and (ii) to fix the rules for typing and equality judgments on F -objects.

The above concepts are related by **theory morphisms**, and theories and morphism form a category. Morphisms subsume all structural relations between theories including logic translations, denotational semantics, interpretation functions, imports, instantiation, inheritance, functors, implementations, and models.

For theories S and T , a morphism $m : S \rightarrow T$ maps every S -symbol to a T -object, which induces a homomorphic translation $m(-)$ of all S -objects to T -objects. MMT guarantees that $m(-)$ preserves the MMT judgments, i.e., $\vdash_S o : o'$ implies $\vdash_T m(o) : m(o')$ and accordingly for equality. This includes the preservation of provability as a special case, which allows moving theorems between theories along morphisms.

Example Consider the representation of various theories in the diagram of Fig. 1. Here the logical framework LF is used as a foundation to define first and higher-order logic. LF contains mostly declarations without types but with notations such as the two declarations

$$\begin{array}{l} \text{type} \\ \text{arrow } \# A_1 \rightarrow A_2 \end{array}$$

which declare a symbol **type** for the collection of types and a binary symbol **arrow** for forming function types $o_1 \rightarrow o_2$. The dotted morphisms represent the meta-relation where one theory is imported to be used as the meta-language in which another theory is defined. For example, FOL uses the symbols and notations of its meta-theory LF in declarations such as




FIGURE 1. Diagram used in the example

```

o      : type
i      : type
equal  : i → i → o # A1 ≐ A2
forall : (i → o) → o # ∀A1

```

which declare symbols

- `o` for the type of formulas,
- `i` for the type of terms,
- `equal` for the binary equality predicate between terms (using curried function application),
- `forall` for the universal quantifier (using higher-order abstract syntax).

FOL in turn is used as the meta-theory for algebraic theories, here the theories of semigroups and rings. For example, `Semigroup` uses the symbols of FOL to declare

```

comp  : i → i → i # A1 ∘ A2
assoc : ded (∀[x : i]∀[y : i]∀[z : i](x ∘ y) ∘ z ≐ x ∘ (y ∘ z))

```

The definition of the theory `Ring` uses the MMT module system to obtain a Bourbaki-style definition that maximizes reuse. The morphisms `add` and `mult` arise from the two imports of the theory `Semigroup` that yield the two different copies of the theory of semigroups present in the theory `Ring`. For example, the associativity axiom is only declared once in `Semigroup` but imported twice into `Ring` and available as `add/assoc` and `mult/assoc`.¹

The morphism `Trans` is a logic translation from first to higher-order logic. And the morphism `Sem` formalizes the semantics of higher-order logic in ZF set theory. As the meta-theory of ZF, we use LF^{\leq} , an extension of `LF` that adds subtyping. The composition `Sem ∘ Trans` yields the semantics of first-order logic in set theory. The morphism `Int` represents the integers as a model of `Ring` defined in set theory, and the compositions `Int ∘ add` and `Int ∘ mult` are the additive and multiplicative semigroup of the integers.

All theories can be moved along morphisms via pushouts. For example, all developments over `LF` can be moved to LF^{\leq} . This has the crucial benefit

¹ Of course, in practical formalizations we use multiple intermediate theories, e.g., for monoids and groups, to fully benefit from modularity.

that we do not lose any work when migrating to a new meta-logic. Similarly, **Semigroup** and **Ring** can be moved to **HOL** by pushout along **Trans** or to **ZF** by pushout along **Sem** \circ **Trans**.

Dia	$::= (Thy \mid Mor)^*$	Diagram
Thy	$::= c[: o] = \{Dec^*\}$	Theory declaration
Mor	$::= c : o \rightarrow o = \{Ass^*\}$	Morphism declaration
Dec	$::= c[: o][= o][\#N]$	Symbol declaration
Ass	$::= c := o$	Assignment to symbol
o	$::= c \mid x \mid Str^c \mid c((x[: o])^*; o^*)$	Object
N	$::= (A_{Int} \mid V_{Int} \mid Str)^*$	Notation
c	$::= URIs$	Identifiers
Str	$::= Unicode\ strings$	Literals or delimiters
Int	$::= integers$	Argument positions

Formal Grammar. The above grammar contains all the essentials of the MMT language. A **diagram** is a list of theory and morphism declarations. The declaration of a **theory** t is of the form $t[: M] = \{\dots, c[: o][= o'][\#N], \dots\}$ where M is the optional meta-theory.

The declaration of a **morphism** m from S to T is of the form $m : S \rightarrow T = \{\dots, c := o, \dots\}$ such that every S -symbol c is mapped to a T -object o .

MMT uses only 4 productions for **T -objects**. 3 of these yields the basic leaves of syntax trees: references to symbols c that are visible to T , references to previously bound variables x , and literals s^c of type c with string representation s (e.g., -1.5^{float} or abc^{string}). The remaining production $c(x_1[: o_1], \dots, x_m[: o_m]; a_1, \dots, a_n)$ subsumes the various ways of forming complex objects: It applies a symbol c , binding some variables x_i , to some arguments a_j . For example, **and**($; F, G$) represents the conjunction of F and G , and **forall**($x; F$) represents the universal quantification over x in F . We obtain the usual notations $F \wedge G$ and $\forall x.F$ if we declare the symbols with notations such as **and** $\#A_1 \wedge A_2$ and **forall** $\#\forall V_1.A_2$.

Specific Foundations. Individual foundations arise as fragments of MMT: A foundation singles out the well-formed theories and objects. We refer to [40] for the details and only sketch the key idea.

MMT provides a minimal inference system for the typing and equality judgments that fixes only the foundation-independent rules, e.g.,

- the well-formedness of theories relative to the well-formed of objects,
- the typing rules for the leaves of the syntax tree, e.g., we have $\vdash_T c : o$ whenever $c : o$ is declared in T ,
- the equality rules for equivalence, congruence, and α -equality.

Further rules are added by the foundations. For example, we can declare a symbol o for the type of well-formed formulas and add a typing-rule for **and** that derives $\vdash_T \text{and}(\ ; F, G) : o$ from $\vdash_T F : o$ and $\vdash_T G : o$.

The foundations only restrict attention to fragments of MMT and do not change the MMT grammar. This is crucial because many definitions can be stated and many theorems proved foundation-independently once and for

all relative to the MMT grammar. This is in contrast to foundation-specific approaches, which have to repeat these steps tediously, often less elegantly, every single time. Examples of such foundation-independent results include

- the category of theories and morphisms and universal constructions such as pushouts,
- the homomorphic extension $m(-)$ and the proof that morphisms preserve judgments,
- the notion of free/bound variables and the definition of capture-avoiding substitution,
- the use of notations to relate concrete and abstract syntax,
- the management of change through differencing and patching,
- the modular development of large theories.

As an example, we consider the module system.

Module System. The foundation-independent module system of MMT provides a very simple and expressive syntax for forming large theories from small components, where the semantics of theory formation is independent of the chosen foundation. We distinguish two ways of forming complex theories and morphisms:

(1) *Declaration-based formation* uses special declarations that elaborate into sets of symbol declarations. Most importantly, the declaration **include** T elaborates into the set of all declarations of T . We can use that to define the theory **Semigroup** more modularly as

$$\begin{aligned} \text{Magma} &: \text{FOL} = \{\text{comp} : \mathbf{i} \rightarrow \mathbf{i} \rightarrow \mathbf{i}\} \\ \text{Semigroup} &: \text{FOL} = \{ \\ &\quad \mathbf{include} \text{ Magma} \\ &\quad \mathbf{assoc} : \mathbf{ded} (\forall [x : \mathbf{i}] \forall [y : \mathbf{i}] \forall [z : \mathbf{i}] (x \circ y) \circ z \doteq x \circ (y \circ z)) \\ &\} \end{aligned}$$

The special declaration $q : T = \{\dots, c := o, \dots\}$ declares an instance q of T . It elaborates into a fresh copy of T where all T -symbols are qualified by q ; moreover, some T -symbols c can be substituted with objects o . This is used in the theory **Ring** to create to different imports of the same theory:

$$\begin{aligned} \text{Ring} &: \text{FOL} = \{ \\ &\quad \mathbf{add} : \text{Semigroup} = \{\} \\ &\quad \mathbf{mult} : \text{Semigroup} = \{\} \\ &\quad \dots \\ &\} \end{aligned}$$

The details can be found in [43].

(2) *Object-based formation* uses symbols to form objects that denote anonymous complex theories or morphisms. This allows forming infinitely many theories and morphisms without adding new declarations.

For example, we declare a unary symbol **identity** $\# \text{id}(A_1)$ and add inference rules that make $\text{id}(T)$ the identity morphism $T \rightarrow T$. Similarly, we declare

a symbol $\text{comp} \# A_1 \circ A_2$ that maps two morphisms to their composition.² We already used this symbol in the example above to obtain the complex morphism $\text{Sem} \circ \text{Trans}$. [42] discusses further constructions in more detail, most importantly pushouts.

3. Foundation-Independent Logic

In MMT, we can give concise foundation-independent definitions of the central concepts regarding logics. This has been described in detail in [40], and we will only give examples here.

Meta-Logic. Meta-logics are represented as foundations in MMT. For example, the theory for LF is shown on the right. It introduces one symbol for each primitive concept along with notations for it. For example, the abstract syntax for a λ -abstraction is $\text{lambda}(x : A; t)$, and the concrete syntax is $[x : A]t$.

Additionally, LF declares one symbol for each rule that is added to MMT's inference system. These are the usual rules for a dependently-typed λ -calculus such as

```

LF = {
kind
type      : kind
Pi        # { V1 } A2
lambda    # [ V1 ] A2
apply     # A1 A2
arrow     # A1 → A2
inferPi
inferlambda
inferapply
funcExten
beta
}

```

$$\frac{\Gamma, x : A \vdash_{\Sigma} t : B}{\Gamma \vdash_{\Sigma} [x : A]t : \{x : A\}B} \text{infer}_{\text{lambda}}$$

Of course, these rules cannot themselves be declared formally.³ Nonetheless, they are declared as symbols of the MMT theory LF and thus subject to the module system. Therefore, e.g., $\text{infer}_{\text{lambda}}$ may be used to derive a typing judgment about T -objects only if T imports LF. Moreover, we can build further meta-logics by importing LF and adding rules for, e.g., rewriting or polymorphism.

Syntax and Proof Theory. For the special case of using LF as a meta-logic, a logic **syntax** is any MMT theory that includes LF and the declarations $\text{o} : \text{type}$ and $\text{ded} : \text{o} \rightarrow \text{type}$. For example, a theory PL for propositional logic could add the declarations on the left and a theory FOL for first-order logic could import PL and add the ones on the right (where we omit the usual notations):

² MMT identifies a symbol by the pair of its name and the name of its containing theory. Therefore, this symbol comp does not clash with the one of the same name from the theory **Semigroup**.

³ The simplest reasonable meta-logic in which we can formalize these rules is LF itself. So we have to supply these rules extra-linguistically to get off the ground. Once a sufficiently expressive meta-logic is represented in this way, we never have to add inference rules again.

\top	:	\circ			
\perp	:	\circ	\mathbf{i}	:	\mathbf{type}
\neg	:	$\circ \rightarrow \circ$	\doteq	:	$\mathbf{i} \rightarrow \mathbf{i} \rightarrow \circ$
\wedge	:	$\circ \rightarrow \circ \rightarrow \circ$	\forall	:	$(\mathbf{i} \rightarrow \circ) \rightarrow \circ$
\vee	:	$\circ \rightarrow \circ \rightarrow \circ$	\exists	:	$(\mathbf{i} \rightarrow \circ) \rightarrow \circ$
\Rightarrow	:	$\circ \rightarrow \circ \rightarrow \circ$			

Given a logic syntax L , an L -**theory** is a theory that extends L with additional declarations. For example, the FOL-theory of semigroups adds

comp	:	$i \rightarrow i \rightarrow i$	$\#A_1 \circ A_2$
assoc	:	$\mathbf{ded} (\forall[x : \mathbf{i}]\forall[y : \mathbf{i}]\forall[z : \mathbf{i}](x \circ y) \circ z \doteq x \circ (y \circ z))$	

A **sentence** over an L -theory T is any object F such that $\vdash_T F : \circ$.

The **proof theory** of a logic is defined as a theory that extends the syntax. For example, some proof rules for a natural deduction calculus for propositional logic are declared as:

$\vee I_l$:	$\{A : \circ\}\{B : \circ\}\mathbf{ded} A \rightarrow \mathbf{ded} [A \vee B]$
$\vee I_r$:	$\{A : \circ\}\{B : \circ\}\mathbf{ded} B \rightarrow \mathbf{ded} [A \vee B]$
$\vee E$:	$\{A : \circ\}\{B : \circ\}\{C : \circ\}\mathbf{ded} [A \vee B] \rightarrow$ $(\mathbf{ded} A \rightarrow \mathbf{ded} C) \rightarrow (\mathbf{ded} B \rightarrow \mathbf{ded} C) \rightarrow \mathbf{ded} C$

Finally, a **proof** of the T -sentence F under the assumptions F_1, \dots, F_n is any object such that $x_1 : \mathbf{ded} F_1, \dots, x_n : \mathbf{ded} F_n \vdash_T p : \mathbf{ded} F$.

Model Theory and Logic Translations. There is an intuitive similarity between model theory and logic translations: Both are inductive translations from one formalism into another: Model theory translates the syntax and proofs into the semantic domain; logic translations translates syntax and proofs to some other logic. In MMT level, we can capture this similarity formally: Both are represented as theory morphisms.


In particular, we represent the semantic domain, in which models are defined, as a theory itself. Traditional logic assumes an implicit mathematical domain such as axiomatic set theory. But modern approaches in computer science often use other, precisely defined domains such as higher-order logic, constructive type theory, or even programming languages. Therefore, MMT allows choosing different semantic domains. Comprehensive examples of foundations and model theory are given in [20, 24].

For example, the set theoretical model theory of HOL is represented as a morphism $\mathbf{Sem} : \mathbf{HOL} \rightarrow \mathbf{ZF}$, where \mathbf{ZF} formalizes axiomatic set theory. \mathbf{Sem} maps every logical symbol to its interpretation—this corresponds to the cases in the inductive definition of the interpretation function. The different ways of extending \mathbf{Sem} to the non-logical symbols declared in FOL-theories T correspond to the different possible T -models. Thus, models are represented as theory morphisms as well.

Finally we can show once and for all and for any meta-logic that every logic defined in MMT induces an institution. This ties together the concrete

representations in logical frameworks and the abstract representations as institutions. We obtain an according result for logic translations, which are represented as theory morphisms $L \rightarrow L'$ such as from FOL to HOL.

The LATIN project [7] leveraged the MMT module system and the above conceptualization to systematically develop a comprehensive library of logics from small reusable components. The whole library includes ≈ 1000 modules, a small fragment is visualized below. All theories use LF as their meta-theory. The left part relates, e.g., propositional, first-order, higher-order, modal, and description logics as well as set theories via imports and translation. The middle part shows how propositional logic PL is build up from individual features for PL, for e.g., negation and conjunction. The right part shows how the module for conjunction consists of formalizations of syntax, model theory, and proof theory.



4. Foundation-Independent Algorithms

The MMT system provides foundation-independent implementations of the typical algorithms needed to obtain tool support for a logic. These are obtained by systematically differentiating between the *foundation-independent* and the *foundation-specific* aspects of existing solutions, and implementing the former in MMT.

Each implementation is relative to a set of rules via which the foundation-specific aspects are supplied. This yields extremely general algorithms that can be easily instantiated for specific foundations. The details of each rule are implemented directly in MMT's underlying programming language (Scala). Practical experience has shown this to be not only feasible but extremely simple and elegant. For example, the foundation-independent parts of MMT consist of $\approx 30,000$ lines of Scala code. But, e.g., the LF plugin provides only 10 simple rules of a few hundred lines of Scala code. Of course, for logics defined using a meta-logic like LF, no such rules have to be provided at all since they are induced by the ones of LF.

Moreover, rules are declared as regular MMT symbols (as we did for the LF in Sect. 3), and the implementations use exactly those rules that are imported into the respective context. That makes it very easy to recombine and extend sets of rules to implement new languages. For example, it took a single rule of 10 lines of code to add shallow polymorphism to LF, and a

single rule schema to add rewriting. This is in contrast to foundation-specific implementations, where similar extensions can require whole PhD theses or may be infeasible altogether.

Below we sketch MMT’s solutions for parsing and type-checking as examples and discuss applications to deduction and computation.

Parsing and Presenting. The MMT algorithms for parsing and presenting are completely foundation-independent.⁴ They use the notations visible to T to parse/present T -objects. Notably, because the grammar for MMT objects is so simple, they are easier to design and implement than their foundation-specific counterparts. Therefore, the foundation-independent MMT algorithms are actually stronger than those of many state-of-the-art logic tools.

As indicated in Sect. 3, MMT notations subsume not only the usual fixity-based notations such as $A_1 \rightarrow A_2$ for the infix notation of `arrow` but also complex notations for binders such as $[V_1]A_2$ for `lambda`, which binds one variable and then takes one argument. But MMT also supports additional advanced features, each only present in very few foundation-specific systems:

- Sequence arguments and sequences of bound variables. For LF, we can actually use notations such as $[V_1, \dots]A_2$ and $(A_1 * \dots) \rightarrow A_2$. The former expresses that `lambda` binds a comma-separated list of variables; the latter expresses that `arrow` takes a star-separated list of arguments.
- Implicit arguments. Notations may declare some arguments to be implicit. These are omitted in concrete syntax, and parsing/type checking must infer them from the context.
- 2-dimensionality. Notations can declare over-/under-/sub-/superscripts as well as fraction-style notations. These are used for presentation in 2-dimensional output formats such as HTML.

Typing. An important advantage of MMT’s grammar for objects is that it can express both human-written and machine-checked syntax trees in the same format. These two are often very different, e.g., when human-written syntax omits inferable subobjects. For example, given the concrete syntax $[x]x \wedge x$, parsing yields the syntax tree `lambda(x; $\wedge(; x, x)$)`. Now type checking this tree infers the type of x and returns `lambda(x : \circ ; $\wedge(; x, x)$)`. Foundation-specific implementations often use a separate intermediate representation format for the former. Because MMT uses the same representation format for both, it can react more flexibly when type checking is not needed or fails.

More generally, the MMT type-checker takes an MMT judgment such as $U \vdash_T o : o'$ where U declares some variables for unknown subobjects. Then it solves for the unique substitution σ such that $\vdash_T o[\sigma] : o'[\sigma]$. This problem is undecidable for most foundations and partial solutions are very difficult to obtain. MMT’s implementation elegantly separates the layers: a foundation-independent core algorithm handles much of the difficulty of the problem; it is customized by supplying a set of rules that handle all foundation-specific aspects and that are comparatively easy to implement. These rules are declared

⁴ Technically, foundation-specific rules may have to be provided for lexing literals. The details can be found in [41].

in MMT theories and thus part of the module system. For example, the rules for LF are the symbols `inferPi`, `inferlambda`, etc. indicated in the definition of LF in Sect. 3. Despite the high generality of MMT’s foundation-independent solution, there are only a few foundation-specific solutions (e.g., the one in Coq [9]) that are substantially stronger.

As an example, we give the implementation of the rule `inferlambda`, which is supplied by the LF plugin:

```
object InferLambda extends TypeInferenceRule(LF.lambda) {
  def apply(solver:Solver,o:Object,context:Context):Option[Object]={
    o match {
      case LF.lambda(x,a,t) =>
        solver.check(Typing(context, a, LF.type))
        solver.inferType(t, context ++ (x oftype a)).map{
          b => LF.Pi(x,a,b)
        }
      case _ => None
    }
  }
}
```

Here `TypeInferenceRule` signals when the core algorithm should apply the rule: during type inference, i.e., when Γ and o are given and o' such that $\Gamma \vdash_{\text{LF}} o : o'$ is needed. Its constructor argument `LF.lambda` makes the applicability more precise: The rule is applicable to `lambda`-objects, i.e., whenever o is of the form `lambda(x : a; t)`. When applied, it receives the object o and the current context Γ and returns the inferred type if possible. It also receives the current instance of the core algorithm `solver`, which maintains the state of the algorithm, in particular the unknown subobjects U . `solver` also offers callbacks for discharging the premises of the rule. In this case, two premises are discharged: The rule checks that $\Gamma \vdash_{\text{LF}} a : \text{type}$, and infers b such that $\Gamma, x : A \vdash_{\text{LF}} t : b$. If the latter succeeds, `Pi(x : a; b)` is returned.

Note how foundation-independence yields a clear separation of concerns. The foundation-specific core—e.g., the rule `inferlambda`—is supplied by the developer of the LF plugin. But the general aspects of type-checking—e.g., constraint propagation and error reporting—are handled foundation-independently by the `solver`. Similarly, the module system remains transparent: The rules are not aware of the modular structure of the current theory. This is in contrast to foundation-specific systems where module system and type checking often have to interact in non-trivial ways.

MMT implements objects as an inductive data type with four constructors for symbols, variables, literals, and complex objects. One might expect that this foundation-independent representation of objects makes it awkward to implement foundation-specific rules. For example, for LF, one would prefer an inductive data type with one constructor each for `type`, `lambda`, `Pi`, and `apply`.

MMT makes this possible: It uses the notations to generate constructor/destructor abbreviations (using Scala’s extractor patterns [13]). Thus, e.g.,

LF-rules can be defined as if there were an LF-specific inductive data type. For example, `LF.Pi(x, a, b)` abbreviates the internal representation of $\text{Pi}(x : a; b)$. Notably, these abbreviations are also available during pattern-matching: `case LF.lambda(x, a, t)` matches any MMT object of the form `lambda(x : a; t)`.

Deduction and Computation. The semantics of logics and related languages can be defined either deductively or computationally. Both paradigms are well-suited for foundation-independent implementation, and the integration of typing, deduction, and computation is an open research question. MMT can state this question foundation-independently, which makes it a good framework for exploring solutions.

In MMT, deduction means to find some object p such that $\vdash_T p : F$, in which case p proves F . Computation means to find some simple object o' such that $\vdash_T o \equiv o'$, in which case o evaluates to o' . MMT provides promising but still very basic foundation-independent implementations for both. In the long run, we can use MMT to build a foundation-independent theorem prover to integrate specialized tools for individual logics. Or we can formalize the semantics of programming languages in MMT in order to reason about programs.

Currently, MMT's foundation-independent theorem prover implements, e.g., the search space, structure sharing, and backtracking. Similarly, its foundation-independent computation engine implements the congruence closure and rewriting. In both cases, the individual proving/computation steps are delegated to the set of rules visible to T . Both algorithms are transparently integrated with typing.

Notably, the LF plugin defines only three deduction rules and already yields a simple theorem prover for any logic defined in LF.

Computation rules for specific theories can be provided by giving models whose semantic domain is a programming language [41]. That allows integrating arbitrary computation with logics.

However, contrary to parsing and type-checking, deduction and computation often require foundation-specific optimizations such as search strategies and decision procedures. Here MMT's foundation-independent implementations are still too weak to compete with foundation-specific ones. But there is no indication they cannot be improved dramatically in future work.

5. Foundation-Independent Knowledge Management

Developing knowledge management services and applications is usually too expensive for individual logic communities, especially if it requires optimization for large scale use cases. Indeed, even the logics with the strongest tool support fare badly on knowledge management. Fortunately, most of this support can be obtained completely foundation-independently—we do not even have to supply any foundation-specific rules as we did in Sect. 4. Thus, maybe surprisingly, foundation-independence helps solve problems in general at relative ease that have proved very hard in each foundation-specific instance.

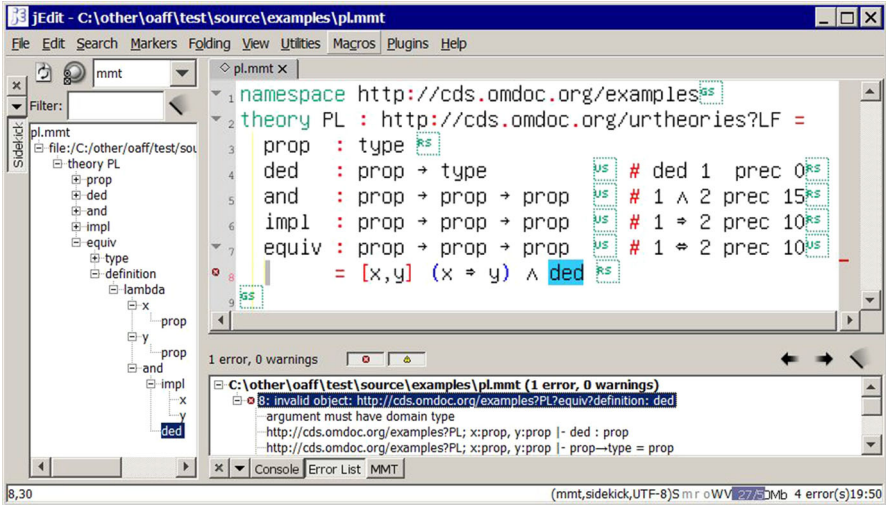


FIGURE 2. Screenshot of the MMT IDE

User Interfaces. MMT provides two foundation-independent user interfaces that go beyond the state-of-the-art of all but very few foundation-specific solutions. It uses jEdit, a mature full-fledged text editor, as the host system for a foundation-independent IDE. And it uses web browsers as the host system for a library browsing environment.

The user interfaces are described in detail in [39], and we only list some of the advanced features. Both include hyperlinking of symbols to their declaration and tooltips that dynamically infer the type of the selected subobject. The IDE displays the list of errors and the abstract syntax tree of a source file as shown in Fig. 2. Both are cross-referenced with the respective source locations. Moreover, the IDE provides context-sensitive auto-completion, which uses the available proving rules to suggest operators that can return an object of the required type. This already yields a basic interactive theorem prover. The web browser interface includes 2-dimensional presentations, e.g., for proof trees, and the folding and hiding of subobjects. It also allows dynamically displaying additional information such as SVG graphs of the modular structure of theories.

Moreover, both interfaces are highly scalable. For example, the IDE employs change management to recheck a declaration only when it was changed or affected by a change. And the multi-threaded HTTP server loads and unloads theories into main memory dynamically so that it can serve very large libraries.

Services. The simplicity of the MMT syntax makes it easy to develop advanced foundation-independent services. We mention only some examples here. [37] defines a **query language** that allows retrieving sets of declarations based on relational (RDF-style) and tree-based (XQuery-style) criteria. Queries may

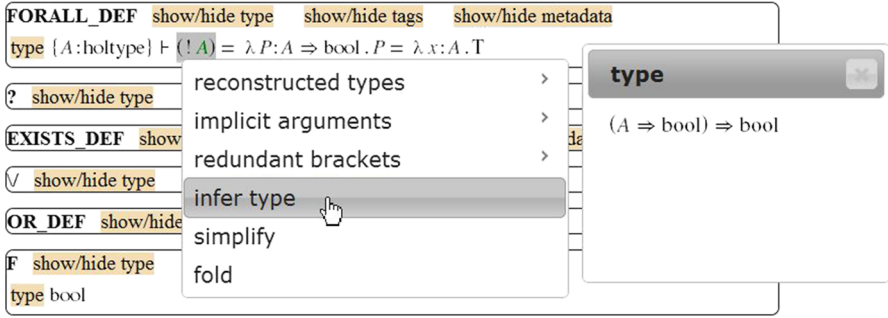


FIGURE 3. Screenshot of the library browser

refer both to the MMT concepts and to user-annotated metadata. MathWebSearch [30] is a massively optimized **substitution tree index** of objects. It performs unification-based search over extremely large libraries almost-instantaneously. [25] develops **change management** support for MMT. It creates differences and patches that only include those changes in an MMT file that are semantically relevant.

All services are independent and exposed through high-level interfaces so that they can be easily reused when building MMT-based applications such as the user interfaces mentioned above. Of particular importance is the MathHub system [22]. Based on git, Drupal, and MMT, it uses the above services in a highly scalable project management and archiving solution for logic libraries.

The viability and strength of these approaches has been demonstrated by instantiating the above services for several major logics such as for Mizar in [23] and for HOL Light in [26]. For example, Fig. 3 shows the dynamic inference in the web-based library browser: Browsing the HOL Light library, the user selected the subexpression `!A` (universal quantification at type A) and called type inference, which returned $(A \Rightarrow \text{bool}) \Rightarrow \text{bool}$. This requires identifying the selected subexpression, inferring its type, and rendering the result using notations—completely foundation-independently apart from using the type inference rules of LF. Foundation-specific systems, on the other hand, are usually designed in such a way that it would be very difficult to offer such an interactive behavior in a web browser at all.

6. Conclusion

The success of logic in the future depends on the solution of one major problem: the proliferation of different logics suffering from incompatible foundations and imperfect and expensive (if any) tool support. These logics and tools are competing instead of collaborating, thus creating massive duplication of work and unexploited synergies. Moreover, new logics are designed much faster than tool support can be developed, e.g., in the area of modal logic. This inefficient

allocation of resources must be overcome for scaling up applications of logic in the future.

Over several decades, three mostly disjoint research communities in logic have independently recognized this problem, and each has developed a major solution: logical frameworks, institutions, and markup languages. All three solutions can be seen as steps towards foundation-independence, where conceptualizations, theorems, and tool support are obtained uniformly for an arbitrary logic.

But these solutions have been developed separately, and each can only solve some aspects of the problem. Logical frameworks are too restrictive and ignore model theoretical aspects. Institutions are too abstract and lack proof theoretical tool support. And markup languages do not formalize the semantics of logics.

The present author has picked up these ideas and coherently re-invented them in a novel framework: the foundation-independent MMT language and tool.

We have described the existing results, which show that MMT allows obtaining simple and powerful results that apply to any logic. Notably, these results range from deep theoretical results to large scale implementations. Within MMT, new logics can be defined extremely easily, and mature, scalable implementations can be obtained at extremely low cost.

The vast majority of logic-related problems can be studied within MMT—in extremely general settings and with tight connections to both theoretical foundations and practical applications. The MMT language is very simple and extensible, and the MMT tool is open-source, well-documented, and systematically designed to be extensible. Thus, it provides a powerful universal framework for the future of logic.

Acknowledgements

The general ideas behind foundation-independence predate the author's research and were already part of his Ph.D. topic as posed by Michael Kohlhase. The MMT language was developed in collaboration with Kohlhase and has incorporated ideas by many members of his research group at Jacobs University Bremen. The MMT tool includes contributions from many group members and has recently become part of a larger software suite that is developed as a group effort.

References

- [1] Anonymous: The QED Manifesto. In: Bundy, A. (ed.) *Automated Deduction*, pp. 238–251. Springer, New York (1994)
- [2] Asperti, A., Sacerdoti Coen, C., Tassi, E., Zacchiroli, S.: Crafting a proof assistant. In: Altenkirch, T., McBride, C. (eds.) *Types for Proofs and Programs*, pp. 18–32. Springer, Berlin, Heidelberg (2007)

- [3] Ausbrooks, R., Buswell, S., Carlisle, D., Dalmas, S., Devitt, S., Diaz, A., Froumentin, M., Hunter, R., Ion, P., Kohlhase, M., Miner, R., Poppelier, N., Smith, B., Soiffer, N., Sutor, R., Watt, S.: Mathematical Markup Language (MathML) Version 2.0 (2nd edition), (2003). See <http://www.w3.org/TR/MathML2>
- [4] Boespflug, M., Carbonneaux, Q., Hermant, O.: The $\lambda\Pi$ -calculus modulo as a universal proof language. In: Pichardie, D., Weber, T. (eds.) Proceedings of PxTP2012: Proof Exchange for Theorem Proving, pp. 28–43 (2012)
- [5] Buswell, S., Caprotti, O., Carlisle, D., Dewar, M., Gaetano, M., Kohlhase, M.: The Open Math Standard, Version 2.0. Technical report, The Open Math Society (2004). See <http://www.openmath.org/standard/om20>
- [6] Church, A.: A formulation of the simple theory of types. *J. Symbol. L.* **5**(1), 56–68 (1940)
- [7] Codescu, M., Horozal, F., Kohlhase, M., Mossakowski, T., Rabe, F.: Project Abstract: logic atlas and integrator (LATIN). In: Davenport, J., Farmer, W., Urban, J., Rabe, F. (eds.) *Intelligent Computer Mathematics*, pp. 289–291. Springer, Berlin, Heidelberg (2011)
- [8] Constable, R., Allen, S., Bromley, H., Cleaveland, W., Cremer, J., Harper, R., Howe, D., Knoblock, T., Mendler, N., Panangaden, P., Sasaki, J., Smith, S.: *Implementing Mathematics with the Nuprl Development System*. Prentice-Hall, Upper Saddle River (1986)
- [9] Coq Development Team. *The Coq Proof Assistant: Reference Manual*. Technical report, INRIA (2015)
- [10] Cousineau, D., Dowek, G.: Embedding pure type systems in the lambda-pi-calculus modulo. In: Della Rocca, S. R. (ed.) *Typed Lambda Calculi and Applications*, pp. 102–117. Springer, Berlin, Heidelberg (2007)
- [11] Curry, H., Feys, R.: *Combinatory Logic*. North-Holland, Amsterdam (1958)
- [12] Diaconescu, R.: *Institution-independent Model Theory*. Birkhäuser, Geneva (2008)
- [13] Emir, B., Odersky, M., Williams, J.: Matching objects with patterns. In: Ernst, E. (ed.) *European Conference on Object-Oriented Programming*, pp. 273–298. Springer, New York (2007)
- [14] Goguen, J., Burstall, R.: Institutions: abstract model theory for specification and programming. *J. Assoc. Comput. Mach.* **39**(1), 95–146 (1992)
- [15] Gordon, M.: HOL: a proof generating system for higher-order logic. In: Birtwistle, G., Subrahmanyam, P. (eds.) *VLSI Specification, Verification and Synthesis*, pp. 73–128. Kluwer-Academic Publishers, Berlin (1988)
- [16] Hales, T., Adams, M., Bauer, G., Tat Dang, D., Harrison, J., Le Hoang, T., Kaliszzyk, C., Magron, V., McLaughlin, S., Tat Nguyen, T., Quang Nguyen, T., Nipkow, T., Obua, S., Pleso, J., Rute, J., Solovyev, A., Thi Ta, A., Nam Tran, T., Thi Trieu, D., Urban, J., Khac Vu, K., Zumkeller, R.: A formal proof of the Kepler conjecture, (2014). [arxiv:1501.02155](https://arxiv.org/abs/1501.02155)
- [17] Harper, R., Honsell, F., Plotkin, G.: A framework for defining logics. *J. Assoc. Comput. Mach.*, **40**(1), 143–184, (1993)
- [18] Harrison, J.: HOL light: a tutorial introduction. In: *Proceedings of the First International Conference on Formal Methods in Computer-Aided Design*, pp. 265–269. Springer, New York (1996)
- [19] HOL4 Development Team. <http://hol.sourceforge.net/>

- [20] Horozal, F., Rabe, F.: Representing model theory in a type-theoretical logical framework. *Theor. Comput. Sci.* **412**(37), 4919–4945 (2011)
- [21] Howard, W.: The formulas-as-types notion of construction. In: To Curry, H.B.(ed.) *Essays on Combinatory Logic, Lambda-Calculus and Formalism*, pp. 479–490. Academic Press, Waltham (1980)
- [22] Iancu, M., Jucovschi, C., Kohlhase, M., Wiesing, T.: System description: MathHub.info. In: Watt, S., Davenport, J., Sexton, A., Sojka, P., Urban, J. (eds.) *Intelligent Computer Mathematics*, pp. 431–434. Springer, New York (2014)
- [23] Iancu, M., Kohlhase, M., Rabe, F., Urban, J.: The Mizar mathematical library in OMDoc: translation and applications. *J. Auto. Reason.* **50**(2), 191–202 (2013)
- [24] Iancu, M., Rabe, F.: Formalizing foundations of mathematics. *Math. Struct. Comput. Sci.* **21**(4), 883–911 (2011)
- [25] Iancu, M., Rabe, F.: Management of change in declarative languages. In: Campbell, J., Carette, J., Dos Reis, G., Jeurig, J., Sojka, P., Sorge, V., Wenzel, M. (eds.) *Intelligent Computer Mathematics*, pp. 325–340. Springer, New York (2012)
- [26] Kaliszzyk, C., Rabe, F.: Towards knowledge management for HOL light. In: Watt, S., Davenport, J., Sexton, A., Sojka, P., Urban, J. (eds.) *Intelligent Computer Mathematics*, pp. 357–372. Springer, New York (2014)
- [27] Kaufmann, M., Manolios, P., Moore, J.: *Computer-Aided Reasoning: An Approach*. Kluwer Academic Publishers, Berlin (2000)
- [28] Klein, G., Andronick, J., Elphinstone, K., Heiser, G., Cock, D., Derrin, P., Elkaduwe, D., Engelhardt, K., Kolanski, R., Norrish, M., Sewell, T., Tuch, H., Winwood, S.: seL4: formal verification of an operating-system kernel. *Commun. ACM* **53**(6), 107–115 (2010)
- [29] Kohlhase, M.: OMDoc: an open markup format for mathematical documents (Version 1.2). Number 4180 in *Lecture Notes in Artificial Intelligence*. Springer, New York (2006)
- [30] Kohlhase, M., Şucan, I.: A search engine for mathematical formulae. In: Ida, T., Calmet, J., Wang, D. (eds.) *Artificial Intelligence and Symbolic Computation*, pp. 241–253. Springer, New York (2006)
- [31] Mossakowski, T., Maeder, C., Lüttich, K.: The heterogeneous tool set. In: Grumberg, O., Huth, M. (eds.) *Tools and Algorithms for the Construction and Analysis of Systems 2007*. *Lecture Notes in Computer Science*, volume **4424**, pp. 519–522 (2007)
- [32] Nipkow, T., Paulson, L., Wenzel, M.: Isabelle/HOL: a proof assistant for higher-order logic. Springer, New York (2002)
- [33] Owre, S., Rushby, J., Shankar, N.: PVS: a prototype verification system. In: Kapur, D. (ed.) *11th International Conference on Automated Deduction (CADE)*, pp. 748–752. Springer, New York (1992)
- [34] Paulson, L.: Isabelle: A Generic Theorem Prover. *Lecture Notes in Computer Science*. Volume 828, Springer, New York (1994)
- [35] Pfenning, F.: Logical frameworks. In: Robinson, J., Voronkov, A. (eds.) *Handbook of automated reasoning*, pp. 1063–1147. Elsevier, Amsterdam (2001)
- [36] Pfenning, F., Schürmann, C.: System description: Twelf - a meta-logical framework for deductive systems. In: Ganzinger, H. (ed.) *Automated deduction*, pp. 202–206 (1999)

- [37] Rabe, F.: A query language for formal mathematical libraries. In: Campbell, J., Carette, J., Dos Reis, G., Jeurig, J., Sojka, P., Sorge, V., Wenzel, M. (eds.) *Intelligent Computer Mathematics*, pp. 142–157. Springer, New York (2012)
- [38] Rabe, F.: The MMT API: a generic MKM system. In: Carette, J., Aspinall, D., Lange, C., Sojka, P., Windsteiger, W. (eds.) *Intelligent Computer Mathematics*, pp. 339–343. Springer, New York (2013)
- [39] Rabe, F.: A logic-independent IDE. In: Benzmler, C., Woltzenlogel Paleo, B. (eds.) *Workshop on User Interfaces for Theorem Provers*, pages 48–60. Elsevier, Amsterdam (2014)
- [40] Rabe, F.: How to identify, translate, and combine logics? *J. Logic Comput.* (2014). doi:[10.1093/logcom/exu079](https://doi.org/10.1093/logcom/exu079)
- [41] Rabe, F.: Generic literals. In: Kerber, M., Carette, J., Kaliszzyk, C., Rabe, F., Sorge, V. (eds.) *Intelligent Computer Mathematics*, pp. 102–117. Springer, New York (2015)
- [42] Rabe, F.: Theory Expressions: A Survey (2015). http://kwarc.info/frabe/Research/rabe_theoexp_15.pdf
- [43] Rabe, F., Kohlhase, M.: A scalable module system. *Inf. Comput.* **230**(1), 1–54 (2013)
- [44] Trybulec, A., Blair, H.: Computer Assisted Reasoning with MIZAR. In: Joshi, A. (ed.) *Proceedings of the 9th International Joint Conference on Artificial Intelligence*, pp. 26–28. Morgan Kaufmann, Burlington (1985)

Florian Rabe
Jacobs University
Campus Ring 1
28759 Bremen, Germany
e-mail: florian.rabe@gmail.com

Received: April 1, 2015.

Accepted: November 1, 2015.