

Chapter 5

Conceptual Graphs

John F. Sowa

Abstract

A conceptual graph (CG) is a graph representation for logic based on the semantic networks of artificial intelligence and the existential graphs of Charles Sanders Peirce. Several versions of CGs have been designed and implemented over the past thirty years. The simplest are the typeless *core* CGs, which correspond to Peirce's original existential graphs. More common are the *extended* CGs, which are a typed superset of the core. The *research* CGs have explored novel techniques for reasoning, knowledge representation, and natural language semantics. The semantics of the core and extended CGs is defined by a formal mapping to and from the ISO standard for Common Logic, but the research CGs are defined by a variety of formal and informal extensions. This article surveys the notation, applications, and reasoning methods used with CGs and their mapping to and from other versions of logic.

5.1 From Existential Graphs to Conceptual Graphs

During the 1960s, graph-based semantic representations were popular in both theoretical and computational linguistics. At one of the most impressive conferences of the decade, Margaret Masterman [21] introduced a graph-based notation, called a *semantic network*, which included a lattice of concept types; Silvio Ceccato [1] presented *correlational nets*, which were based on 56 different relations, including subtype, instance, part-whole, case relations, kinship relations, and various kinds of attributes; and David Hays [15] presented *dependency graphs*, which formalized the notation developed by the linguist Lucien Tesnière [40]. The early graph notations represented the relational structures underlying natural language semantics, but none of them could express full first-order logic. Woods [42] and McDermott [22] wrote scathing critiques of their logical weaknesses.

In the late 1970s, many graph notations were designed to represent first-order logic or a formally-defined subset [7]. Sowa [32] developed a version of *conceptual graphs* (CGs) as an intermediate language for mapping natural language questions and assertions to a relational database. Fig. 5.1 shows a CG for the sentence *John is going to*




Figure 5.1: CG display form for *John is going to Boston by bus.*

Boston by bus. The rectangles are called *concepts*, and the circles are called *conceptual relations*. An arc pointing toward a circle marks the first *argument* of the relation, and an arc pointing away from a circle marks the last argument. If a relation has only one argument, the arrowhead is omitted. If a relation has more than two arguments, the arrowheads are replaced by integers 1, . . . , n .

The conceptual graph in Fig. 5.1 represents a typed or sorted version of logic. Each of the four concepts has a *type label*, which represents the type of entity the concept refers to: *Person*, *Go*, *Boston*, or *Bus*. Two of the concepts have *names*, which identify the referent: *John* or *Boston*. Each of the three conceptual relations has a type label that represents the type of relation: *agent* (*Agnt*), *destination* (*Dest*), or *instrument* (*Inst*). The CG as a whole indicates that the person *John* is the agent of some instance of going, the city *Boston* is the destination, and a bus is the instrument. Fig. 5.1 can be translated to the following formula:

$$(\exists x)(\exists y)(\text{Go}(x) \wedge \text{Person}(\text{John}) \wedge \text{City}(\text{Boston}) \wedge \text{Bus}(y) \wedge \\ \text{Agnt}(x, \text{John}) \wedge \text{Dest}(x, \text{Boston}) \wedge \text{Inst}(x, y)).$$

As this translation shows, the only logical operators used in Fig. 5.1 are conjunction and the existential quantifier. Those two operators are the most common in translations from natural languages, and many of the early semantic networks could not represent any others.

For his pioneering *Begriffsschrift* (concept writing), Frege [8] adopted a tree notation for representing full first-order logic, using only four operators: assertion (the “turnstile” operator \vdash), negation (a short vertical line), implication (a hook), and the universal quantifier (a cup containing the bound variable). Fig. 5.2 shows the *Begriffsschrift* equivalent of Fig. 5.1, and following is its translation to predicate calculus:

$$\sim(\forall x)(\forall y)(\text{Go}(x) \supset (\text{Person}(\text{John}) \supset (\text{City}(\text{Boston}) \supset \\ (\text{Bus}(y) \supset (\text{Agnt}(x, \text{John}) \supset (\text{Dest}(x, \text{Boston}) \supset \sim \text{Inst}(x, y))))))))$$

Frege’s choice of operators simplified his rules of inference, but they lead to awkward paraphrases: *It is false that for every x and y , if x is an instance of going then if *John* is a person then if *Boston* is a city then if y is a bus then if the agent of x is *John* then if the destination of x is *Boston* then the instrument of x is not y .*

Unlike Frege, who rejected Boolean algebra, Peirce developed the algebraic notation for first-order logic as a generalization of the Boolean operators. Since Boole treated disjunction as logical addition and conjunction as logical multiplication,




Figure 5.2: Frege's Begriffsschrift for *John is going to Boston by bus*.

Peirce [24] represented the existential quantifier by Σ for repeated disjunction and the universal quantifier by Π for repeated conjunction. In the notation of Peirce [25], Fig. 5.1 could be represented

$$\Sigma_x \Sigma_y (\text{Go}(x) \bullet \text{Person}(\text{John}) \bullet \text{City}(\text{Boston}) \bullet \text{Bus}(y) \bullet \text{Agnt}(x, \text{John}) \bullet \text{Dest}(x, \text{Boston}) \bullet \text{Inst}(x, y)).$$

Peano adopted Peirce's notation, but he invented new symbols because he wanted to mix mathematical and logical symbols in the same formulas. Meanwhile, Peirce began to experiment with *relational graphs* for representing logic, as in Fig. 5.3. In that graph, an existential quantifier is represented by a *line of identity*, and conjunction is the default Boolean operator. Since Peirce's graphs did not distinguish proper names, the monadic predicates *isJohn* and *isBoston* may be used to represent names. Following is the algebraic notation for Fig. 5.3:

$$\Sigma_x \Sigma_y \Sigma_z \Sigma_w (\text{Go}(x) \bullet \text{Person}(y) \bullet \text{isJohn}(y) \bullet \text{City}(z) \bullet \text{isBoston}(z) \bullet \text{Bus}(w) \bullet \text{Agnt}(x, y) \bullet \text{Dest}(x, z) \bullet \text{Inst}(x, w)).$$

Peirce experimented with various graphic methods for representing the other operators of his algebraic notation, but like the AI researchers of the 1960s, he could not find a good way of expressing the scope of quantifiers and negation. In 1897, however, he discovered a simple, but brilliant innovation for his new version of *existential graphs* (EGs): an oval enclosure for showing scope [27]. The default operator for an oval with no other marking is negation, but any metalevel relation can be linked to the oval. Sowa [33] adopted Peirce's convention for CGs, but with rectangles instead of ovals: rectangles nest better than ovals; and more importantly, each context box can be interpreted as a concept box that contains a nested CG. A nest of two negations




Figure 5.3: Peirce's relational graph for *John is going to Boston by bus*.




Figure 5.4: EG and CG for *If a farmer owns a donkey, then he beats it.*

indicates an implication, as in Fig. 5.4, which shows an EG and a CG for the sentence *If a farmer owns a donkey, then he beats it.*

As Fig. 5.4 illustrates, the primary difference between EGs and CGs is the interpretation of the links: in EGs, each line of identity represents an existentially quantified variable, which is attached to the relations; in CGs, the concept boxes represent existential quantifiers, and the arcs merely link relation nodes to their arguments. Another difference is that the CG type labels become monadic relations in EGs. Unlike EGs, in which an unmarked oval represents negation, the symbol \sim marks a negated CG context. Both the EG and the CG could be represented by the following formula:

$$\sim(\exists x)(\exists y)(\text{Farmer}(x) \wedge \text{Donkey}(y) \wedge \text{Owns}(x, y) \wedge \sim \text{Beats}(x, y)).$$

In order to preserve the correct scope of quantifiers, the implication operator \supset cannot be used to represent the English *if-then* construction unless the existential quantifiers are moved to the front and converted to universals:

$$(\forall x)(\forall y)((\text{Farmer}(x) \wedge \text{Donkey}(y) \wedge \text{Owns}(x, y)) \supset \text{Beats}(x, y)).$$

In English, this formula could be read *For every x and y, if x is a farmer who owns a donkey y, then x beats y.* The unusual nature of this paraphrase led Kamp [18] to develop *discourse representation structures* (DRSs) whose logical structure is isomorphic to Peirce’s existential graphs (Fig. 5.5).

Kamp’s primitives are the same as Peirce’s: the default quantifier is the existential, and the default Boolean operator is conjunction; negation is represented by a context box, and implication is represented by two contexts. As Fig. 5.5 illustrates, the nesting




Figure 5.5: EG and DRS for *If a farmer owns a donkey, then he beats it.*

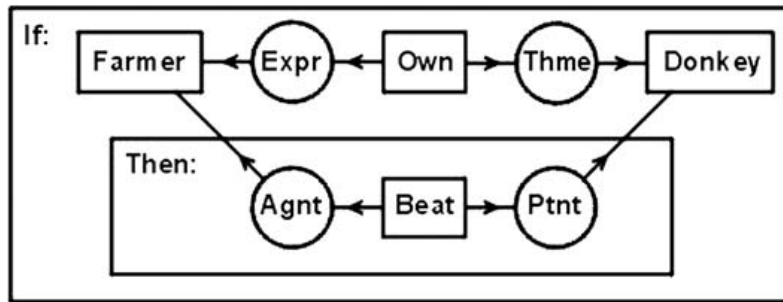


Figure 5.6: CG with case relations shown explicitly.

of Peirce's contexts allows the quantifiers in the antecedent of an implication to include the consequent within their scope. Although Kamp connected his boxes with an arrow, he made exactly the same assumption about the scope of quantifiers. Kamp and Reyle [19] went much further than Peirce in analyzing discourse and formulating the rules for interpreting anaphoric references, but any rule stated in terms of the DRS notation can also be applied to the EG or CG notation.

The CG in Fig. 5.4 represents the verbs *owns* and *beats* as dyadic relations. That was the choice of relations selected by Kamp, and it can also be used with the EG or CG notation. Peirce, however, noted that the event or state expressed by a verb is also an entity that could be referenced by a quantified variable. That point was independently rediscovered by linguists, computational linguists, and philosophers such as Davidson [6]. The CG in Fig. 5.6 shows a representation that treats events and states as entities linked to their participants by *case relations* or *thematic roles*.

The type labels *If* and *Then* in Fig. 5.6 are defined as synonyms for negated contexts. The state of owning is linked to its participants by the relations experiencer (*Expr*) and theme (*Thme*), and the act of beating by the relations agent (*Agnt*) and patient (*Ptnt*). Following is the equivalent in typed predicate calculus:

$$\begin{aligned} &\sim(\exists x:\text{Farmer})(\exists y:\text{Own})(\exists z:\text{Donkey})(\text{Expr}(y, x) \wedge (\text{Thme}(y, z) \wedge \\ &\sim(\exists w:\text{Beat})(\text{Agnt}(w, x) \wedge \text{Ptnt}(w, z))). \end{aligned}$$

The model-theoretic semantics for the EGs and CGs shown in this section is specified in the ISO standard for Common Logic (CL) [17]. Section 5.2 of this article briefly describes the CL project, the CL model theory, and the mapping of the CL abstract syntax to and from the Conceptual Graph Interchange Format (CGIF), a linear notation that represents every semantic feature of the graphs. Section 5.3 presents the *canonical formation rules* for CGs and their use with Peirce's rules of inference for full FOL. Section 5.4 presents the use of CGs for representing propositions, situations, and metalevel reasoning. Section 5.5 discusses research issues that have inspired a variety of formal and informal extensions to the conceptual graph theory and notation.

5.2 Common Logic

Common Logic (CL) evolved from two projects to develop parallel ANSI standards for conceptual graphs and the Knowledge Interchange Format [9]. Eventually, those

two projects were merged into a single ISO project to develop a common abstract syntax and model-theoretic foundation for a family of logic-based notations [17]. Hayes and Menzel [13] defined a very general model theory for CL, which Hayes and McBride [12] used to define the semantics for the languages RDF(S) and OWL. In addition to the abstract syntax and model theory, the CL standard specifies three concrete dialects that are capable of expressing the full CL semantics: the Common Logic Interchange Format (CLIF), the Conceptual Graph Interchange Format (CGIF), and the XML-based notation for CL (XCL). RDF and OWL can also be considered dialects that express subsets of the CL semantics: any statement in RDF or OWL can be translated to CLIF, CGIF, or XCL, but only a subset can be translated back to RDF or OWL.

The CL syntax allows quantifiers to range over functions and relations, but CL retains a first-order style of model theory and proof theory. To support a higher-order syntax, but without the computational complexity of higher-order semantics, the CL model theory uses a single domain D that includes individuals, functions, and relations. The option of limiting the domain of quantification to a single set was suggested by Quine [29] and used in various theorem provers that allow quantifiers to range over relations [3].

Conceptual graphs had been a typed version of logic since the first publication in 1976, but Peirce's untyped existential graphs are sufficiently general to express the full CL semantics. Therefore, two versions of the Conceptual Graph Interchange Format are defined in the ISO standard:

1. **Core CGIF.** A typeless version of logic that expresses the full CL semantics. This dialect corresponds to Peirce's existential graphs: its only primitives are conjunction, negation, and the existential quantifier. It does permit quantifiers to range over relations, but Peirce also experimented with that option for EGs.
2. **Extended CGIF.** An upward compatible extension of the core, which adds a universal quantifier; type labels for restricting the range of quantifiers; Boolean contexts with type labels *If*, *Then*, *Either*, *Or*, *Equivalence*, and *Iff*; and the option of importing external text into any CGIF text.

Although extended CGIF is a typed language, it is not *strongly typed*, because type labels are used only to restrict the range of quantifiers. Instead of causing a syntax error, as in the strongly typed logic Z [16], a type mismatch in CGIF just causes the subexpression in which the mismatch occurs to be false. If a typed sentence in Z is translated to CGIF, it will have the same truth value in both languages, but a type mismatch, such as the following, is handled differently in each:

```
~[ [Elephant: 23] ]
```

This CGIF sentence, which is syntactically correct and semantically true, says that 23 is not an elephant. If translated to Z , however, the type mismatch would cause a syntax error. The more lenient method of handling types is necessary for representing sentences derived from other languages, both natural and artificial. RDF and OWL, for example, can be translated to CGIF and CLIF, but not to Z .

The conceptual graph in Fig. 5.1, which represents the sentence *John is going to Boston by bus*, can be written in the following form in extended CGIF:

```
[Go *x] [Person: John] [City: Boston] [Bus *y]
(Agnt ?x John) (Dest ?x Boston) (Inst ?x ?y)
```

In CGIF, concepts are marked by square brackets, and conceptual relations are marked by parentheses. A character string prefixed with an asterisk, such as **x*, marks a *defining node*, which may be referenced by the same string prefixed with a question mark, *?x*. These strings, which are called *name sequences* in Common Logic, represent *coreference labels* in CGIF and variables in other versions of logic. Following is the equivalent in CLIF:

```
(exists ((x Go) (y Bus))
  (and (Person John) (city Boston)
    (Agnt x John) (Dest x Boston) (Inst x y) ))
```

In the CL standard, extended CGIF is defined by a translation to core CGIF, which is defined by a translation to the CL abstract syntax. Following is the untyped core CGIF and the corresponding CLIF for the above examples:

```
[*x] [*y]
(Go ?x) (Person John) (City Boston) (Bus ?y)
(Agnt ?x John) (Dest ?x Boston) (Inst ?x ?y)

(exists (x y)
  (and (Go x) (Person John) (city Boston) (Bus y)
    (Agnt x John) (Dest x Boston) (Inst x y) ))
```

In core CGIF, the most common use for concept nodes is to represent existential quantifiers. A node such as *[*x]* corresponds to an EG line of identity, such as the one attached to the relation *Go* in Fig. 5.3. It is permissible to write names in a concept node such as *[: John]*, but in most cases, such nodes are unnecessary because names can also be written in relation nodes. A concept node may contain more than one name or coreference label, such as *[: John ?z]*. In EGs, that node corresponds to a *ligature* that links two lines of identity; in CLIF, it corresponds to an equality: *(= John z)*.

Although CGIF and CLIF look similar, there are several fundamental differences:

1. Since CGIF is a serialized representation of a graph, labels such as *x* or *y* represent connections between nodes in CGIF, but variables in CLIF or predicate calculus.
2. Since the nodes of a graph have no inherent ordering, a CGIF sentence is an unordered list of nodes. Unless grouped by context brackets, the list may be permuted without affecting the semantics.
3. The CLIF operator *and* does not occur in CGIF because the conjunction of nodes within any context is implicit. Omitting the conjunction operator in CGIF tends to reduce the number of parentheses.

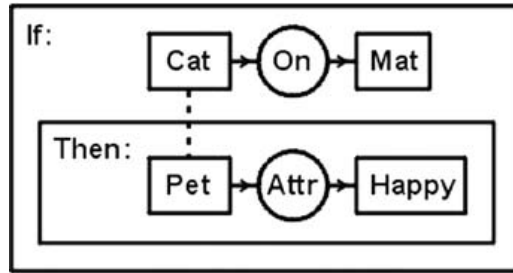


Figure 5.7: CG display form for *If a cat is on a mat, then it is a happy pet.*

4. Since CGIF labels show connections of nodes, they may be omitted when they are not needed. One way to reduce the number of labels is to move concept nodes inside the parentheses of relation nodes:

```
[Go *x]
  (Agnt ?x [Person: John])
  (Dest ?x [City: Boston])
  (Inst ?x [Bus])
```

When written in this way, CGIF looks like a frame notation. It is, however, much more general than frames, since it can represent the full semantics of CL.

As another example, Fig. 5.7 shows a CG for the sentence *If a cat is on a mat, then it is a happy pet.* The dotted line that connects the concept [Cat] to the concept [Pet], which is called a *coreference link*, indicates that they both refer to the same entity. The Attr relation indicates that the cat, also called a pet, has an attribute, which is an instance of happiness.

The coreference link in Fig. 5.7 is shown in CGIF by the defining label *x in the concept [Cat : *x] and the bound label ?x in the concept [Pet : ?x]. Following is the extended CGIF and its translation to core CGIF:

```
[If: [Cat *x] [Mat *y] (On ?x ?y)
  [Then: [Pet ?x] [Happy *z] (Attr ?x ?z) ]]

~[ [*x] [*y] (Cat ?x) (Mat ?y) (On ?x ?y)
  ~[ (Pet ?x) [*z] (Happy ?z) (Attr ?x ?z) ]]
```

In CGs, functions are represented by conceptual relations called *actors*. Fig. 5.8 is the CG display form for the following equation written in ordinary algebraic notation:

$$y = (x + 7) / \text{sqrt}(7)$$

The three functions in this equation would be represented by three actors, which are shown in Fig. 5.8 as diamond-shaped nodes with the type labels Add, Sqrt, and Divide. The concept nodes contain the input and output values of the actors. The two empty concept nodes contain the output values of Add and Sqrt.

In CGIF, actors are represented as relations with two kinds of arcs: a sequence of *input arcs* and a sequence of *output arcs*, which are separated by a vertical bar:

```
[Number: *x] [Number: *y] [Number: 7]
(Add ?x 7 | [*u]) (Sqrt 7 | [*v]) (Divide ?u ?v | ?y)
```





Figure 5.8: CL functions represented by actor nodes.

In the display form, the input arcs of `Add` and `Divide` are numbered 1 and 2 to indicate the order in which the arcs are written in CGIF. Following is the corresponding CLIF:

```
(exists ((x Number) (y Number))
  (and (Number 7) (= y (Divide (Add x 7) (Sqrt 7)))))
```

No CLIF variables are needed to represent the coreference labels `*u` and `*v` since the functional notation used in CLIF shows the connections directly.

CLIF only permits functions to have a single output, but extended CGIF allows actors to have multiple outputs. The following actor of type `IntegerDivide` has two inputs: an integer `x` and an integer `7`. It also has two outputs: a quotient `u` and a remainder `v`.

```
(IntegerDivide [Integer: *x] [Integer: 7] | [*u] [*v])
```

When this actor is translated to core CGIF or CLIF, the vertical bar is removed, and the actor becomes an ordinary relation with four arguments; the distinction between inputs and outputs is lost. In order to assert the constraint that the last two arguments are functionally dependent on the first two arguments, the following CGIF sentence asserts that there exist two functions, identified by the coreference labels `Quotient` and `Remainder`, which for every combination of input and output values are logically equivalent to an actor of type `IntegerDivide` with the same input and output values:

```
[Function: *Quotient] [Function: *Remainder]
[[@every*x1] [@every*x2] [@every*x3] [@every*x4]
[Equiv: [Iff: (IntegerDivide ?x1 ?x2 | ?x3 ?x4)]
  [Iff: (#?Quotient ?x1 ?x2 | ?x3)
    (#?Remainder ?x1 ?x2 | ?x4)]]]
```

Each line of this example illustrates one or more features of CGIF. The first line represents existential quantifiers for two entities of type `Function`. On the second line, the context bracket `[` encloses the concept nodes with universal quantifiers, marked by `@every`, to show that the existential quantifiers for `Quotient` and `Remainder` include the universals within their scope. The equivalence on the last three lines shows that an actor of type `IntegerDivide` is logically equivalent to a conjunction of the quotient and remainder functions. Finally, the symbol `#` on the last two

lines shows that the coreference labels ?Quotient and ?Remainder are being used as type labels. Following is the corresponding CLIF:

```
(exists ((Quotient Function) (Remainder Function))
 (forall (x1 x2 x3 x4)
  (iff (IntegerDivide x1 x2 x3 x4)
   (and (= x3 (Quotient x1 x2)) (= x4 (Remainder x1 x2))))))
```

As another example of the use of quantification over relations, someone might say “Bob and Sue are related”, but not say exactly how they are related. The following sentences in CGIF and CLIF state that there exists some familial relation *r* that relates Bob and Sue:

```
[Relation: *r] (Familial ?r) (#?r Bob Sue)

(exists ((r Relation)) (and (Familial r) (r Bob Sue)))
```

The concept [Relation: *r] states that there exists a relation *r*. The next two relations state that *r* is familial and *r* relates Bob and Sue.

This brief survey has illustrated nearly every major feature of CGIF and CLIF. One important feature that has not been mentioned is the use of *sequence variables* to support relations with a variable number of arguments. Another is the use of comments, which can be placed before, after, or inside any concept or relation node in CGIF. The specifications in the CL standard guarantee that any sentence expressed in any of the three fully conformant dialects—CLIF, CGIF, or XCL—can be translated to any of the others in a logically equivalent form. Although the translation will preserve the semantics, it is not guaranteed to preserve all syntactic details: a sentence translated from one dialect to another and then back to the first will be logically equivalent to the original, but some subexpressions might be reordered or replaced by semantic equivalents.

In general, Common Logic is a superset of many different logic-based languages and notations, including the traditional predicate-calculus notation for first-order logic. But since various languages have been designed and implemented at widely separated times and places, that generalization must be qualified with different caveats for each case:

1. **Semantic Web Languages.** The draft CL standard supports the URIs defined by the W3C as valid CL name sequences, and it allows text stored on the web to be imported into CLIF, CGIF, or XCL documents. The tools that import the text could, if necessary, translate one dialect to another at import time. Since the semantics for RDF(S) and OWL was designed as a subset of the CL model theory, those languages can be translated to any fully conformant CL dialect [11].
2. **Z Specification Notation.** The Z model theory is a subset of the CL model theory, but the syntax of Z enforces strong type checking, and it does not permit quantifiers to range over functions and relations. Therefore, Z statements can be translated to CL, but only those statements that originally came from Z are guaranteed to be translatable back to Z.

3. **Unified Modeling Language (UML).** Although the UML diagrams and notations are loosely based on logic, they have no formal specification in any version of logic. The best hope for providing a reliable foundation for UML would be to implement tools that translate UML to CL. If done properly, such tools could define a *de facto* standard for UML semantics.
4. **Logic-Programming Languages.** Well-behaved languages that support classical negation can be translated to CL while preserving the semantics. Languages based on negation as failure, such as Prolog, could be translated to CL, but with the usual caveats about ways of working around the discrepancies.
5. **SQL Database Language.** The WHERE clause in SQL queries and constraints can state an arbitrary FOL expression, but problems arise with the treatment of null values in the database and with differences between the open-world and closed-world assumptions. To avoid the nonlogical features of SQL, CL can be mapped to and from the Datalog language, which supports the Horn-clause subset of FOL and has a direct mapping to the SQL operations.

Most people have strong attachments to their favorite syntactic features. The goal of the Common Logic project is to provide a very general semantics that enables interoperability at the semantic level despite the inevitable syntactic differences. CL has demonstrated that such seemingly diverse notations as conceptual graphs, predicate calculus, and the languages of the Semantic Web can be treated as dialects with a common semantic foundation. An extension of CL called IKL, which is discussed in Section 5.5, can support an even wider range of logics.

5.3 Reasoning with Graphs

Graphs have some advantages over linear notations in both human factors and computational efficiency. As Figs. 5.1–5.8 illustrate, graphs show relationships at a glance that are harder to see in linear notations, including CGIF and CLIF. Graphs also have a highly regular structure that can simplify many algorithms for reasoning, searching, indexing, and pattern matching. Yet AI research has largely ignored the structural properties of graphs, and some of the most advanced research on representing, indexing, and manipulating graphs has been done in organic chemistry. With his BS degree in chemistry, Peirce was the first to recognize the similarity between chemical graphs and logical graphs. He wanted to represent the “atoms and molecules of logic” in his existential graphs, and he used the word *valence* for the number of arguments of a relation. By applying algorithms for chemical graphs to conceptual graphs, Levinson and Ellis [20] implemented the first type hierarchy that could support retrieval and classification in logarithmic time. More recent research on chemical graphs has been used in algorithms for computing semantic distance between CGs. Those techniques have enabled analogy finding in logarithmic time, instead of the polynomial-time computations of the older methods [37].

The six *canonical formation rules* [34] are examples of graph-based operators that focus on the semantics. Combinations of these rules, called *projection* and *maximal join*, perform larger semantic operations, such as answering a question or comparing

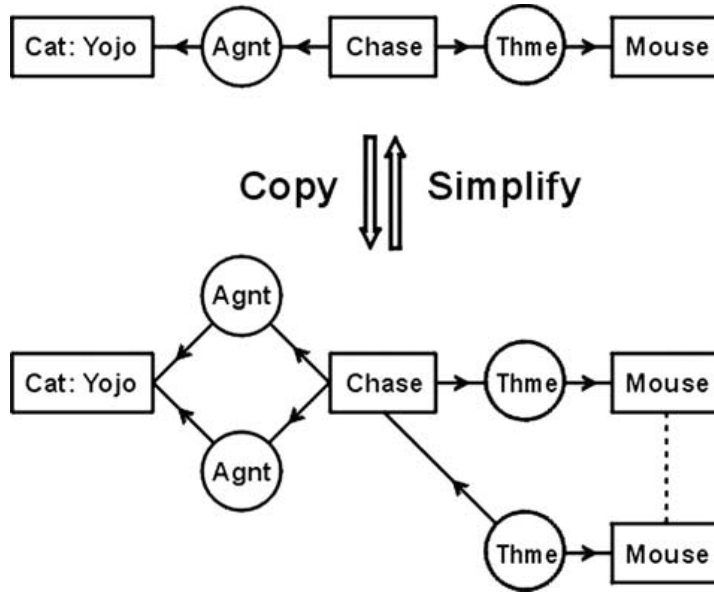


Figure 5.9: Copy and simplify rules.

the relevance of different alternatives. Each rule has one of three possible effects on the logical relationship between a starting graph u and the resulting graph v :

1. **Equivalence.** *Copy* and *simplify* are equivalence rules, which generate a graph v that is logically equivalent to the original: $u \supset v$ and $v \supset u$. Equivalent graphs are true in exactly the same models.
2. **Specialization.** *Join* and *restrict* are specialization rules, which generate a graph v that implies the original: $v \supset u$. Specialization rules monotonically decrease the set of models in which the result is true.
3. **Generalization.** *Detach* and *unrestrict* are generalization rules, which generate a graph v that is implied by the original: $u \supset v$. Generalization rules monotonically increase the set of models in which the result is true.

Each rule has an inverse rule that undoes any change caused by the other. The inverse of copy is simplify, the inverse of restrict is unrestrict, and the inverse of join is detach. These rules are fundamentally graphical: they are easier to show than to describe. Figures 5.9 to 5.11 illustrate these rules with *simple graphs*, which use only conjunction and existential quantifiers. When rules for handling negation are added, they form a complete proof procedure for first-order logic with equality.

The CG at the top of Fig. 5.9 represents the sentence *The cat Yojo is chasing a mouse*. The down arrow represents two applications of the copy rule. One application copies the Agnt relation, and a second copies the subgraph $\rightarrow (\text{Thme}) \rightarrow [\text{Mouse}]$. The coreference link connecting the two [Mouse] concepts indicates that both concepts refer to the same individual. The up arrow represents two applications of the simplify rule, which performs the inverse operations of erasing redundant copies. Following are the CGIF sentences for both graphs:




Figure 5.10: Restrict and unrestrict rules.

[Cat: Yojo] [Chase: *x] [Mouse: *y]
 (Agent ?x Yojo) (Thme ?x ?y)

[Cat: Yojo] [Chase: *x] [Mouse: *y] [Mouse: ?y]
 (Agent ?x Yojo) (Agent ?x Yojo) (Thme ?x ?y) (Thme ?x ?y)

As the CGIF illustrates, the copy rule makes redundant copies, which are erased by the simplify rule. In effect, the copy rule is $p \supset (p \wedge p)$, and the simplify rule is $(p \wedge p) \supset p$.

The CG at the top of Fig. 5.10 represents the sentence *A cat is chasing an animal*. By two applications of the restrict rule, it is transformed to the CG for *The cat Yojo is chasing a mouse*. In the first step, the concept [Cat], which says that there exists some cat, is *restricted by referent* to the more specific concept [Cat: Yojo], which says that there exists a cat named Yojo. In the second step, the concept [Animal], which says that there exists an animal, is *restricted by type* to a concept of a subtype [Mouse]. The more specialized graph implies the more general one: if the cat Yojo is chasing a mouse, then a cat is chasing an animal.

To show that the bottom graph v of Fig. 5.10 implies the top graph u , let c be a concept of u that is being restricted to a more specialized concept d , and let u be $c \wedge w$, where w is the remaining information in u . By hypothesis, $d \supset c$. Therefore, $(d \wedge w) \supset (c \wedge w)$. Hence, $v \supset u$.

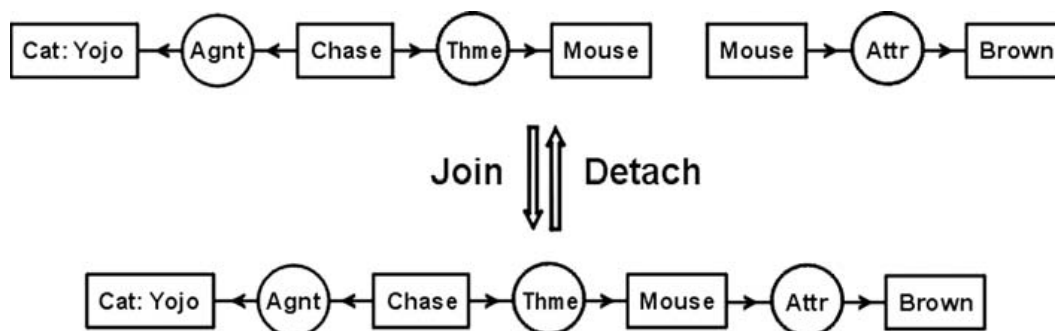


Figure 5.11: Join and detach rules.

At the top of Fig. 5.11 are two CGs for the sentences *Yojo is chasing a mouse* and *A mouse is brown*. The join rule overlays the two identical copies of the concept [Mouse] to form a single CG for the sentence *Yojo is chasing a brown mouse*. The detach rule undoes the join to restore the top graphs. Following are the CGIF sentences that represent the top and bottom graphs of Fig. 5.11:

```
[Cat: Yojo] [Chase: *x] [Mouse: *y] (Agent ?x Yojo)
(Thme ?x ?y) [Mouse: *z] [Brown: *w] (Attr ?z ?w)
```

```
[Cat: Yojo] [Chase: *x] [Mouse: *y] (Agent ?x Yojo)
(Thme ?x ?y) [Brown: *w] (Attr ?y ?w)
```

As the CGIF illustrates, the bottom graph consists of substituting y for every occurrence of z in the top graph and erasing redundant copies. In general, every join assumes an equality of the form $y = z$ and simplifies the result. If q is the equality and u is original pair of graphs at the top, then the bottom graph is equivalent to $q \wedge u$, which implies u . Therefore, the result of join implies the original graphs.

Together, the generalization and equivalence rules are sufficient for a complete proof procedure for the subset of logic whose only operators are conjunction and the existential quantifier. The specialization and equivalence rules can be used in a refutation procedure for a proof by contradiction. To handle full first-order logic, rules for negations must be added. Peirce defined a complete proof procedure for FOL whose rules depend on whether a context is positive (nested in an even number of negations, possibly zero) or negative (nested in an odd number of negations). Those rules are grouped in three pairs, one of which (i) inserts a graph and the other (e) erases a graph. The only axiom is a blank sheet of paper (an empty graph with no nodes); in effect, the blank is a generalization of all other graphs. Following is a restatement of Peirce's rules in terms of specialization and generalization. These same rules apply to both propositional logic and full first-order logic. In FOL, the operation of inserting or erasing a connection between two nodes has the effect of identifying two nodes (i.e., a substitution of a value for a variable) or treating them as possibly distinct.

1. (i) In a negative context, any graph or subgraph (including the blank) may be replaced by any specialization.
(e) In a positive context, any graph or subgraph may be replaced by any generalization (including the blank).
2. (i) Any graph or subgraph in any context c may be copied in the same context c or into any context nested in c . (The only exception is that no graph may be copied directly into itself; however, it is permissible to copy a graph g in the context c and then to copy the copy into the original g .)
(e) Any graph or subgraph that could have been derived by rule 2(i) may be erased. (Whether or not the graph was in fact derived by 2(i) is irrelevant.)
3. (i) A double negation (nest of two negations with nothing between the inner and outer) may be drawn around any graph, subgraph, or set of graphs in any context.
(e) Any double negation in any context may be erased.




Figure 5.12: Proof of Frege's first axiom by Peirce's rules.

This version of the rules was adapted from a tutorial on existential graphs by Peirce [28]. He originally formulated these rules in 1897 and finally published them in 1906, but they are a simplification and generalization of the rules of *natural deduction* by Gentzen [10]. When these rules are applied to CGIF, some adjustments may be needed to rename coreference labels or to convert a bound label to a defining label or vice versa. For example, if a defining node is erased, a bound label may become the new defining label. Such adjustments are not needed in the pure graph notation.

All the axioms and rules of inference for classical FOL, including the rules of the *Principia Mathematica*, natural deduction, and resolution, can be proved in terms of Peirce's rules. As an example, Frege's first axiom, written in Peirce–Peano notation, is $a \supset (b \supset a)$. Fig. 5.12 shows a proof by Peirce's rules. (To improve contrast, positive areas are shown in white, and negative areas are shaded.)

In core CGIF, the propositions a and b could be represented as relations with zero arguments. Following are the five steps of Fig. 5.12 written in core CGIF:

1. By rule 3(i), Insert a double negation around the blank: $\sim[\sim[\]]$.
2. By 3(i), insert a double negation around the previous one:
 $\sim[\sim[\sim[\sim[\]]]]$.
3. By 1(i), insert (a): $\sim[\ (a) \sim[\sim[\sim[\]]]]$.
4. By 2(i), copy (a): $\sim[\ (a) \sim[\sim[\sim[\ (a) \]]]]$.
5. By 1(i), insert (b): $\sim[\ (a) \sim[\sim[\ (b) \sim[\ (a) \]]]]$.

The theorem to be proved contains five symbols, and each step of the proof inserts one symbol into its proper place in the final result. Frege had a total of eight axioms, and the *Principia* had five. All of them could be derived by similarly short proofs.

Frege's two rules of inference, which Whitehead and Russell adopted, were *modus ponens* and *universal instantiation*. Fig. 5.13 is a proof of *modus ponens*, which derives q from a statement p and an implication $p \supset q$:

Following are the four steps of Fig. 5.13 written in core CGIF:

1. Starting graphs: $(p) \sim[\ (p) \sim[\ (q) \]]$.
2. By 2(e), erase the nested copy of (p): $(p) \sim[\ \sim[\ (q) \]]$.



Figure 5.13: Proof of modus ponens.




Figure 5.14: Proof of universal instantiation.

3. By 1(e), erase (p): $\sim[\sim[(q)]]$.
4. By 3(e), erase the double negation: (q) .

Frege's other rule of inference is *universal instantiation*, which allows any term t to be substituted for a universally quantified variable in a statement of the form $(\forall x)P(x)$. In EGs, the term t would be represented by a graph of the form $\text{---}t$, which states that something satisfying the condition t exists, and the universal quantifier corresponds to a negated existential: a line whose outermost part (the existential quantifier) occurs in a negative area. Since a graph has no variables, there is no notion of substitution. Instead, the proof in Fig. 5.14 performs the equivalent operation by connecting the two lines.

The absence of labels on the EG lines simplifies the proofs by eliminating the need to relabel variables in CLIF or coreference links in CGIF. In core CGIF, the first step is the linear version of Fig. 5.14:

1. Starting graphs: $[*x] (t \text{ ?}x) \sim[[*y] \sim[(P \text{ ?}y)]]$.
2. By 2(i), copy $[*x]$ and change the defining label $*x$ to a bound label $?x$ in the copy: $[*x] (t \text{ ?}x) \sim[[?x] [*y] \sim[(P \text{ ?}y)]]$.
3. By 1(i), insert a connection between the two lines. In CGIF, that corresponds to relabeling $*y$ and $?y$ to $?x$ and erasing redundant copies of $[?x]$:
 $[*x] (t \text{ ?}x) \sim[\sim[(P \text{ ?}x)]]$.
4. By 3(e), erase the double negation: $[*x] (t \text{ ?}x) (P \text{ ?}x)$.

With the universal quantifier in extended CGIF, the starting graphs of Fig. 5.14 could be written

$$[*x] (t \text{ ?}x) [(P [\text{@every} *y])].$$

The extra brackets around the last node ensure that the existential quantifier $[*x]$ includes the universal $\text{@every} *y$ within its scope. Then universal instantiation could be used as a one-step derived rule to generate line 4. After @every has been erased, the brackets around the last node are not needed and may be erased.

In the *Principia Mathematica*, Whitehead and Russell proved the following theorem, which Leibniz called the *Praeclarum Theorema* (Splendid Theorem). It is one of the last and most complex theorems in propositional logic in the *Principia*, and it required a total of 43 steps, starting from five nonobvious axiom schemata

$$((p \supset r) \wedge (q \supset s)) \supset ((p \wedge q) \supset (r \wedge s)).$$




Figure 5.15: Proof in 7 steps instead of 43 in the *Principia*.

With Peirce’s rules, this theorem can be proved in just seven steps starting with a blank sheet of paper (Fig. 5.15). Each step of the proof inserts or erases one graph, and the final graph is the statement of the theorem.

After only four steps, the graph looks almost like the desired conclusion, except for a missing copy of s inside the innermost area. Since that area is positive, it is not permissible to insert s directly. Instead, Rule 2(i) copies the graph that represents $q \supset s$. By Rule 2(e), the next step erases an unwanted copy of q . Finally, Rule 3(e) erases a double negation to derive the conclusion.

Unlike Gentzen’s version of natural deduction, which uses a method of making and discharging assumptions, Peirce’s proofs proceed in a straight line from a blank sheet to the conclusion: every step inserts or erases one subgraph in the immediately preceding graph. As Fig. 5.15 illustrates, the first two steps of any proof that starts with a blank must draw a double negation around the blank and insert a graph into the negative area. That graph is usually the entire hypothesis of the theorem to be proved. The remainder of the proof develops the conclusion in the doubly nested blank area. Those two steps are the equivalent of Gentzen’s method of making and discharging an assumption, but in Gentzen’s approach, the two steps may be separated by arbitrarily many intervening steps, and a system of bookkeeping is necessary to keep track of the assumptions. With Peirce’s rules, the second step follows immediately after the first, and no bookkeeping is required.

Most common proofs take about the same number of steps with Peirce’s rules as with Gentzen’s version of natural deduction or his *sequent calculus*. For some kinds of proofs, however, Peirce’s rules can be much faster because of a property that is not shared by any other common proof procedure: the rules depend only on whether an area is positive or negative; the depth of nesting is irrelevant. That property implies the “cut-and-paste theorem” [34], which is proved in terms of Peirce’s rules, but it can be used in any notation for first-order logic:

1. **Theorem.** If a proof $p \vdash q$ is possible on a blank sheet, then in any positive area of a graph or formula where p occurs, q may be substituted for p .
2. **Proof.** Since the area in which p occurs is positive, every step of the proof of q can be carried out in that area. Therefore, it is permissible to “cut out” and “paste in” the steps of the proof from p to q in that area. After q has been

derived, Rule 1(e) can be applied to erase the original p and any remaining steps of the proof other than q .

Dau [4] showed that certain proofs that take advantage of this theorem or the features of Peirce's rules that support it can be orders of magnitude shorter than proofs based on other rules of inference. Conventional rules, for example, can only be applied to the outermost operator. If a graph or formula happens to contain a deeply nested subformula p , those rules cannot replace it with q . Instead, many steps may be needed to bring p to the surface of some formula to which conventional rules can be applied. An example is the *cut-free* version of Gentzen's sequent calculus, in which proofs can sometimes be exponentially longer than proofs in the usual version. With Peirce's rules, the corresponding cut-free proofs are only longer by a polynomial factor.

The canonical formation rules have been implemented in nearly all CG systems, and they have been used in formal logic-based methods, informal case-based reasoning, and various computational methods. A multistep combination, called a *maximal join*, is used to determine the extent of the unifiable overlap between two CGs. In natural language processing, maximal joins can help resolve ambiguities and determine the most likely connections of new information to background knowledge and the previous context of a discourse. Stewart [38] implemented Peirce's rules of inference in a first-order theorem prover for EGs and showed that its performance is comparable to resolution theorem provers. In all reasoning methods, formal and informal, a major part of the time is spent in searching for relevant rules, axioms, or background information. Ongoing research on efficient methods of indexing graphs and selecting the most relevant information has shown great improvement in many cases, but more work is needed to incorporate such indexing into conventional reasoning systems.

5.4 Propositions, Situations, and Metalanguage

Natural languages are highly expressive systems that can state anything that has ever been stated in any formal language or logic. They can even express metalevel statements about themselves, their relationships to other languages, and the truth of any such statements. Such enormous expressive power can easily generate contradictions and paradoxes, such as the statement *This sentence is false*. Most formal languages avoid such paradoxes by imposing restrictions on the expressive power. Common Logic, for example, can represent any sentence in any CL dialect as a quoted string, and it can even specify the syntactic structure of such strings. But CL has no mechanism for treating such strings as CL sentences and relating substrings in them to the corresponding CL names.

Although the paradoxes of logic are expressible in natural language, the most common use of language about language is to talk about the beliefs, desires, and intentions of the speaker and other people. Many versions of logic and knowledge representation languages, including conceptual graphs, have been used to express such language. As an example, the sentence *Tom believes that Mary wants to marry a sailor*, contains three clauses, whose nesting may be marked by brackets:

Tom believes that [Mary wants [to marry a sailor]].

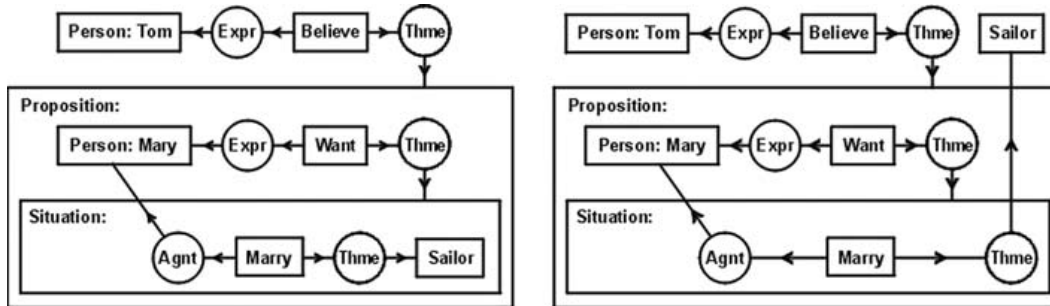


Figure 5.16: Two interpretations of *Tom believes that Mary wants to marry a sailor.*

The outer clause asserts that Tom has a belief, which is expressed by the object of the verb *believe*. Tom's belief is that Mary wants a situation described by the nested infinitive, whose subject is the same person who wants the situation. Each clause makes a comment about the clause or clauses nested in it. References to the individuals mentioned in those clauses may cross context boundaries in various ways, as in the following two interpretations of the original English sentence:

Tom believes that
[there is a sailor whom Mary wants [to marry]].

There is a sailor whom Tom believes that
[Mary wants [to marry]].

The two conceptual graphs in Fig. 5.16 represent the first and third interpretations. In the CG on the left, the existential quantifier for the concept [Sailor] is nested inside the situation that Mary wants. Whether such a sailor actually exists and whether Tom or Mary knows his identity are undetermined. The CG on the right explicitly states that such a sailor exists; the connections of contexts and relations imply that Tom knows him and that Tom believes that Mary also knows him. Another option (not shown) would place the concept [Sailor] inside the context of type *Proposition*; it would leave the sailor's existence undetermined, but it would imply that Tom believes he exists and that Tom believes Mary knows him.

The context boxes illustrated in Figs. 5.4 and 5.6 express negations or operators such as *If* and *Then*, which are defined in terms of negations. However, the contexts of the types *Proposition* and *Situation* in Fig. 5.16 raise new issues of logic and ontology. The CL semantics can represent entities of any type, including propositions and situations, but it has no provision for relating such entities to the internal structure of CL sentences. A more expressive language, called IKL [14], was defined as an upward compatible extension of CL. The IKL semantics introduces entities called *propositions* and a special operator, spelled *that*, which relates IKL sentences to the propositions they express. IKL semantics does not have a built-in type for situations, but it is possible in IKL to make statements that state the existence of entities of type *Situation* and relate them to propositions.

The first step toward translating the CGs in Fig. 5.16 to IKL is to write them in an extended version of CGIF, which allows CGs to be nested inside concept nodes of type *Proposition* or *Situation*. Following is the CGIF for the CG on the left:

```
[Person: Tom] [Believe: *x1] (Expr ?x1 Tom)
(Thme ?x1 [Proposition:
  [Person: Mary] [Want: *x2] (Expr ?x2 Mary)
  (Thme ?x2 [Situation:
    [Marry: *x3] [Sailor: *x4] (Agnt ?x3 Mary)
    (Thme ?x3 ?x4)]))])])])
```

This statement uses the option to move the concept nodes for the types `Proposition` and `Situation` inside the relation nodes of type `Thme`. That option has no semantic significance, but it makes the order of writing the CGIF closer to English word order. A much more important semantic question is the relation between situations and propositions. In the ontology commonly used with CGs, that relation is spelled `Dscr` and called the *description relation*. The last two lines of the CGIF statement above could be rewritten in the following form:

```
(Thme ?x2 [Situation: *s] (Dscr ?s [Proposition:
  [Marry: *x3] [Sailor: *x4] (Agnt ?x3 Mary)
  (Thme ?x3 ?x4)]))
```

The last line is unchanged, but the line before it states that the theme of `x2` is the situation `s` and the description of `s` is the proposition stated on the last line. In effect, every concept of type `Situation` that contains a nested CG is an abbreviation for a situation that is described by a concept of type `Proposition` that has the same nested CG. This expanded CGIF statement can then be translated to IKL (which is based on CLIF syntax with the addition of the operator `that`).

```
(exists ((x1 Believe)) (and (Person Tom) (Expr x1 Tom)
(Thme x1 (that
  (exists ((x2 Want) (s Situation))
    (and (Person Mary) (Expr x2 Mary)
      (Thme x2 s) (Dscr s (that
        (exists ((x3 Marry) (x4 Sailor))
          (and (Agnt x3 Mary)
            (Thme x3 x4) ))))))))))))
```

Each line of the IKL statement expresses the equivalent of the corresponding line in CGIF. Note that every occurrence of `Proposition` in CGIF corresponds to `that` in IKL. The syntax of CLIF or IKL requires more parentheses than CGIF because every occurrence of `(exists` or `(and` requires an extra closing parenthesis at the end.

As these examples illustrate, the operator `that` adds an enormous amount of expressive power, but IKL still has a first-order style of semantics. The proposition nodes in CGs or the `that` operator in IKL introduce abstract entities of type `Proposition`, but propositions are treated as zero-argument relations, which are supported by the semantics of Common Logic. Although language about propositions is a kind of metalanguage, it does not, by itself, go beyond first-order logic. Tarski [39], for example, demonstrated how a stratified series of metalevels, each of which is purely first order, can be used without creating paradoxes or going beyond the semantics of FOL. In effect, Tarski avoided paradoxes by declaring that certain kinds of sentences (those that violate the stratification) do not express propositions in his models. The

IKL model theory has a similar way of avoiding paradoxes: it does not require every model to include a proposition for every possible sentence. For example, the following English sentence, which sounds paradoxical, could be expressed in either IKL or CGIF syntax:

*There exists a proposition p , p is true,
and p is the proposition that p is false.*

Since IKL does not require every sentence to express a proposition in every model, there are permissible IKL models in which this sentence is false simply because no such proposition exists. Therefore, the paradox vanishes because the sentence has a stable, but false, truth value.

5.5 Research Extensions

Over the years, the term *conceptual graph* has been used in a broad sense as any notation that has a large overlap with the notation used in the book by Sowa [33]. That usage has led to a large number of dialects with varying degrees of compatibility. The purpose of a standard is to stabilize a design at a stage where it can provide a fixed, reliable platform for the development of products and applications. A fixed design, however, is an obstacle to innovation in the platform itself, although it is valuable for promoting innovation in applications that use the platform. In order to support fundamental research while providing a stable platform for applications, it is important to distinguish ISO standard CGs, IKL CGs, and research CGs. The first two provide rich and stable platforms for application development, while the third allows research projects to add extensions and modifications, which may be needed for a particular application and which may someday be added to the standard.

Most of the features of the research CGs are required to support natural language semantics. Some of them, such as modal operators, are as old as Aristotle, but they are not in the CL standard because their semantics involves open research issues. Following are the most common research extensions that have been proposed or implemented in various forms over the years:

1. **Contexts.** Peirce's first use for the oval was to negate the graphs nested inside, and that is the only use that is recognized by the CL standard. But Peirce [26] generalized the ovals to context enclosures, which allow relations other than negation to be linked to a context. Most of those features could be defined in terms of the IKL extensions described in Section 5.4, but there is no consensus on any definitions that could be considered for a standard.
2. **Metalanguage.** The basic use of a context enclosure is to quote the nested graphs. That metalevel syntax allows any semantic approach to be defined by axioms that specify how the nested graphs are interpreted. Sowa [35, 36] adapted that approach to a family of *nested graph models* (NGMs), which could be used to formalize the semantics of many kinds of modal and intentional logics. A hierarchy of metalevels with the NGM semantics can express the equivalent of a wide range of modal, temporal, and intentional logics. The most useful NGMs can be represented with the IKL semantics, but the many

variations and their application to natural languages have not yet been fully explored.

3. **Generalized quantifiers.** In addition to the usual quantifiers of *every* and *some*, natural languages support an open-ended number of quantificational expressions, such as *exactly one*, *at least seven*, or *considerably more*. Some of these quantifiers, such as *exactly one cat*, could be represented as [Cat : @1] and defined in terms of the CL standard. Others, such as *at least seven cats*, could be represented [Cat : @≤7] and defined with a version of set theory added to the base logic. But quantifiers such as *considerably more* would require some method of approximate reasoning, such as fuzzy sets or rough sets.
4. **Indexicals.** Peirce observed that every statement in logic requires at least one indexical to fix the referents of its symbols. The basic indexical, which corresponds to the definite article *the*, is represented by the symbol # inside a concept node: [Dog : #] would represent the phrase *the dog*. The pronouns *I*, *you*, and *she* would be represented [Person : #I], [Person : #you], and [Person : #she]. To process indexicals, some linguists propose versions of *dynamic semantics*, in which the model is updated during the discourse. A simpler method is to treat the # symbol as a syntactic marker that indicates a incomplete interpretation of the original sentence. With this approach, the truth value of a CG that contains any occurrences of # is not determined until those markers are replaced by names or coreference labels. This approach supports indexicals in an intermediate representation, but uses a conventional model theory to evaluate the final resolution.
5. **Plural nouns.** Plurals have been represented in CGs by set expressions inside the concept boxes. The concept [Cat : {*}@3] would represent *three cats*, and [Dog : {Lucky, Macula}] would represent *the dogs Lucky and Macula*. Various methods have been proposed for representing distributed and collective plurals and translating them to versions of set theory and mereology. But the representation of plurals is still a research area in linguistics, and it is premature to adopt a standard syntax or semantics.
6. **Procedural attachments.** The CL standard defines actors as purely functional relations, but various implementations have allowed more informal versions, in which the actors represent arbitrary procedures. Some versions have implemented *token passing* algorithms with the symbol ? for a backward-chaining request used in a *query graph*, and with the symbol ! for a forward-chaining trigger that asserts a new value. As examples, the concept [Employee : ?] would ask *which employee*, and the concept [Employee : John!] would assert *the employee John*.

As an example of applied research, one of the largest commercial CG systems is Sonetto [30], which uses extended versions of the earlier algorithms by Levinson and Ellis [20]. A key innovation of Sonetto is its semi-automated methods for extracting ontologies and business rules from unstructured documents. The users who assist Sonetto in the knowledge extraction process are familiar with the subject matter, but they have no training in programming or knowledge engineering. CGIF is the

knowledge representation language for ontologies, rules, and queries. It is also used to manage the schemas of documents and other objects in the system and to represent the rules that translate CGIF to XML and other formats. For the early CG research, see the collections edited by Nagle et al. [23], Way [41], and Chein [2]. More recent research on CGs has been published in the annual proceedings of the International Conferences on Conceptual Structures [5, 31].

Bibliography

- [1] S. Ceccato. *Linguistic Analysis and Programming for Mechanical Translation*. Gordon and Breach, New York, 1961.
- [2] M. Chein, editor, *Revue d'intelligence artificielle*, 10(1), 1996 (Numéro Spécial Graphes Conceptuels).
- [3] W. Chen, M. Kifer, and D.S. Warren. Hilog: A foundation for higher-order logic programming. *Journal of Logic Programming*, 15(3):187–230, 1993.
- [4] F. Dau. Some notes on proofs with Alpha graphs. In [31], pages 172–188, 2006.
- [5] F. Dau, M.-L. Mugnier, and G. Stumme, editors. *Conceptual Structures: Common Semantics for Sharing Knowledge, LNAI*, vol. 3596. Springer, Berlin, 2005.
- [6] D. Davidson. The logical form of action sentences. Reprinted in D. Davidson. *Essays on Actions and Events*. Clarendon Press, Oxford, pages 105–148, 1980.
- [7] N.V. Findler, editor, *Associative Networks: Representation and Use of Knowledge by Computers*. Academic Press, New York, 1979.
- [8] G. Frege, *Begriffsschrift*. 1879; English translation in J. van Heijenoort, editor. *From Frege to Gödel*, pages 1–82. Harvard University Press, Cambridge, MA, 1967.
- [9] M.R. Genesereth and R. Fikes, editors. *Knowledge Interchange Format, Version 3.0 Reference Manual, TR Logic-92-1*. Computer Science Department, Stanford University, 1992.
- [10] G. Gentzen. Untersuchungen über das logische Schließen. In *The Collected Papers of Gerhard Gentzen* (Investigations into logical deduction), pages 68–131. North-Holland Publishing Co., Amsterdam, 1969 (edited and translated by M.E. Szabo).
- [11] P. Hayes. Translating semantic web languages into Common Logic. <http://www.ihmc.us/users/phayes/CL/SW2SCL.html>, 2005.
- [12] P. Hayes and B. McBride. RDF semantics, W3C Technical Report. <http://www.w3.org/TR/rdf-mt/>, 2003.
- [13] P. Hayes and C. Menzel. A semantics for the knowledge interchange format. In *Proc. IJCAI 2001 Workshop on the IEEE Standard Upper Ontology*, Seattle, 2001.
- [14] P. Hayes and C. Menzel. IKL Specification Document. <http://www.ihmc.us/users/phayes/IKL/SPEC/SPEC.html>, 2006.
- [15] D.G. Hays. Dependency theory: a formalism and some observations. *Language*, 40(4):511–525, 1964.
- [16] ISO/IEC, *Z Formal Specification Notation—Syntax, Type System, and Semantics*, IS 13568, International Organisation for Standardisation, 2002.
- [17] ISO/IEC 24707, Information Technology, Common Logic (CL)—A Framework for a family of Logic-Based Languages, International Organisation for Standardisation, Geneva, 2007.

- [18] H. Kamp. A theory of truth and semantic representation. In J.A.G. Groenendijk, T.M.V. Janssen, and M.B.J. Stokhof, editors. *Formal Methods in the Study of Language*, pages 277–322. Mathematical Centre Tracts, Amsterdam, 1981.
- [19] H. Kamp and U. Reyle. *From Discourse to Logic*. Kluwer, Dordrecht, 1993.
- [20] R.A. Levinson and G. Ellis. Multilevel hierarchical retrieval. *Knowledge Based Systems*, 5(3):233–244, 1992.
- [21] M. Masterman. Semantic message detection for machine translation, using an interlingua. In *Proc. 1961 International Conf. on Machine Translation*, pages 438–475, 1961.
- [22] D.V. McDermott. Artificial intelligence meets natural stupidity. *SIGART Newsletter*, 57, April 1976.
- [23] T.E. Nagle, J.A. Nagle, L.L. Gerholz, and P.W. Eklund, editors. *Conceptual Structures: Current Research and Practice*. Ellis Horwood, New York, 1992.
- [24] C.S. Peirce. On the algebra of logic. *American Journal of Mathematics*, 3:15–57, 1880.
- [25] C.S. Peirce. On the algebra of logic. *American Journal of Mathematics*, 7:180–202, 1885.
- [26] C.S. Peirce. Reasoning and the Logic of Things. In K.L. Ketner, editor. *The Cambridge Conferences Lectures of 1898*. Harvard University Press, Cambridge, MA, 1992.
- [27] C.S. Peirce. Manuscripts on existential graphs. In *Collected Papers of Charles Sanders Peirce*, vol. 4, pages 320–410. Harvard University Press, Cambridge, MA, 1906.
- [28] C.S. Peirce. Manuscript 514, with commentary by J.F. Sowa, <http://www.jfsowa.com/peirce/ms514.htm>, 1909.
- [29] W.V. Orman Quine. Reduction to a dyadic predicate. *J. Symbolic Logic*, 19, 1954. Reprinted In W.V. Quine, editor. *Selected Logic Papers*, enlarged edition, pages 224–226. Harvard University Press, Cambridge, MA, 1995.
- [30] Q. Sarraf and G. Ellis. Business rules in retail: The Tesco.com story. *Business Rules Journal*, 7(6), 2006, <http://www.brcommunity.com/a2006/n014.html>.
- [31] H. Schärfe, P. Hitzler, and P. Øhrstrom, editors. *Conceptual Structures: Inspiration and Application*, *LNAI*, vol. 4068. Springer, Berlin, 2006.
- [32] J.F. Sowa. Conceptual graphs for a database interface. *IBM Journal of Research and Development*, 20(4):336–357, 1976.
- [33] J.F. Sowa. *Conceptual Structures: Information Processing in Mind and Machine*. Addison-Wesley, Reading, MA, 1984.
- [34] J.F. Sowa. *Knowledge Representation: Logical, Philosophical, and Computational Foundations*. Brooks/Cole Publishing Co., Pacific Grove, CA, 2000.
- [35] J.F. Sowa. Laws, facts, and contexts: Foundations for multimodal reasoning. In V.F. Hendricks, K.F. Jørgensen, and S.A. Pedersen, editors. *Knowledge Contributors*, pages 145–184. Kluwer Academic Publishers, Dordrecht, 2003.
- [36] J.F. Sowa. Worlds, models, and descriptions. *Studia Logica, Ways of Worlds II*, 84(2):323–360, 2006 (Special Issue).
- [37] J.F. Sowa and A.K. Majumdar. Analogical reasoning. In A. de Moor, W. Lex, and B. Ganter, editors. *Conceptual Structures for Knowledge Creation and Communication*, *LNAI*, vol. 2746, pages 16–36. Springer-Verlag, Berlin, 2003.
- [38] J. Stewart. Theorem proving using existential graphs. MS thesis, Computer and Information Science, University of California at Santa Cruz, 1996.

- [39] A. Tarski. The concept of truth in formalized languages. In A. Tarski, editor. *Logic, Semantics, Metamathematics*, 2nd edition, pages 152–278. Hackett Publishing Co., Indianapolis, 1933.
- [40] L. Tesnière. *Éléments de Syntaxe structurale*, 2nd edition. Librairie C. Klincksieck, Paris, 1965.
- [41] E.C. Way, editor, *Journal of Experimental and Theoretical Artificial Intelligence (JETAI)*, 4(2), 1992 (Special Issue on Conceptual Graphs).
- [42] W.A. Woods. What's in a link: foundations for semantic networks. In D.G. Bobrow and A. Collins, editors. *Representation and Understanding*, pages 35–82. Academic Press, New York, 1975.

This page intentionally left blank