Developing and Empirically Evaluating Robust Explanation Generators: The KNIGHT Experiments*

James C. Lester
Department of Computer Science
North Carolina State University
Raleigh, NC 27695-8206
(919) 515-7135
fax: (919) 515-7896
lester@adm.csc.ncsu.edu

Bruce W. Porter
Department of Computer Sciences
University of Texas at Austin
Austin, Texas 78712-1188
(512) 471-9578
fax: (512) 471-8885
porter@cs.utexas.edu

December 1, 1995

Abstract

To explain complex phenomena, an explanation system must be able to select information from a formal representation of domain knowledge, organize the selected information into multi-sentential discourse plans, and realize the discourse plans in text. Although recent years have witnessed significant progress in the development of sophisticated computational mechanisms for explanation, empirical results have been limited. This paper reports on a seven year effort to empirically study explanation generation from semantically rich, large-scale knowledge bases. In particular, it describes KNIGHT, a robust explanation system that constructs multi-sentential and multi-paragraph explanations from the *Biology Knowledge Base*, a large-scale knowledge base in the domain of botanical anatomy, physiology, and development. We introduce the *Two Panel* evaluation methodology and describe how KNIGHT's performance was assessed with this methodology in the most extensive empirical evaluation conducted on an explanation system. In this evaluation, KNIGHT scored within "half a grade" of domain experts, and its performance exceeded that of one of the domain experts.

Keywords: Explanation, natural language generation, discourse planning, large-scale knowledge bases.

*Support for this research is provided by a grant from the National Science Foundation (IRI-9120310), a contract from the Air Force Office of Scientific Research (F49620-93-1-0239), and donations from the Digital Equipment Corporation. This research was conducted at the University of Texas at Austin.

1 Introduction

In the course of their daily affairs, scientists explain complex phenomena—both to one another and to lay people—in a manner that facilitates clear communication. Similarly, physicians, lawyers, and teachers are equally facile at generating explanations in their respective areas of expertise. In an effort to computationalize this critical ability, research in natural language generation has addressed a broad range of issues in automatically constructing text from formal representations of domain knowledge. Research on text planning [42, 52, 39, 24, 38] has developed techniques for determining the content and organization of many genres, and explanation generation [44, 8, 47] in particular has been the subject of intense investigation. In addition to exploring a panorama of application domains, the explanation community has begun to assemble these myriad designs into a coherent framework. As a result, we have begun to see a crystalization of the major components, as well as detailed analyses of their roles in explanation [62, 63].

Despite this success, empirical results in explanation generation are limited. Although techniques for developing and evaluating robust explanation generation should yield results that are more conclusive than those produced by prototype, "proof-of-concept" systems, with only a few notable exceptions [27, 23, 8, 46, 54], most work has adopted a research methodology in which a proof-of-concept system is constructed and its operation is analyzed on a few examples. While isolating one or a small number of problems enables researchers to consider particular issues in detail, it is difficult to gauge the scalability and robustness of a proposed approach.

A critical factor contributing to the dearth of empirical results is the absence of semantically rich, large-scale knowledge bases. Knowledge bases housing tens of thousands of different concepts and hundreds of different relations could furnish ample raw materials for empirical study, but no work in explanation generation has been conducted or empirically evaluated in the context of these knowledge bases.

To empirically study explanation generation from semantically rich, large-scale knowledge bases, we undertook a seven year experiment. First, our domain experts (one employed full-time) constructed the Biology Knowledge Base [53], a very large structure representing more than 180,000 facts about botanical anatomy, physiology, and development. Second, we designed, implemented, and empirically evaluated KNIGHT [29], a robust explanation system that extracts information from the Biology Knowledge Base, organizes it, and realizes it in multi-sentential and multi-paragraph expository explanations of complex biological phenomena. Third, we developed a novel evaluation methodology for gauging the effectiveness of explanation systems and employed this methodology to evaluate KNIGHT.

This paper describes the lessons learned during the course of the "KNIGHT experiments." In the spirit of EDGE [8] and PAULINE [16], which synthesize work in interactive explanation systems and generational pragmatics, respectively, KNIGHT addresses a broad range of issues, all in the context of semantically rich, large-scale knowledge bases:

- *Robust Knowledge-Base Access:* KNIGHT exploits a library of robust knowledge-base access methods that insulate discourse planners from the idiosyncracies and errors in knowledge bases. These "view construction" methods selectively extract coherent packets of propositions about the structure and function of objects, the changes made to objects by processes, and the temporal attributes and temporal decompositions of processes.

- *Discourse-Knowledge Engineering:* Discourse-knowledge engineers, i.e., knowledge engineers who encode discourse knowledge, should be able to inspect and easily modify discourse-planning specifications for rapid iterative refinement. The *Explanation Design Package* (EDP) formalism is a convenient, schema-like [41, 52] programming language for text planning. Because the EDP formalism is a hybrid of the declarative and procedural paradigms, discourse-knowledge engineers can easily understand EDPs, modify them, and use them to represent new discourse knowledge. EDPs have been used by KNIGHT to generate hundreds of expository explanations of biological objects and processes.

- *Explanation Planning:* KNIGHT employs a robust explanation planner that selects EDPs and applies them to invoke knowledge-base accessors. The explanation planner considers the desired length of explanations and the relative importance of sub-topics as it constructs explanation plans encoding content and organization.

- *Functional Realization:* KNIGHT's functional realization system [7] is built on top of a unification-based surface generator with a large systemic grammar [15, 16].

To assess KNIGHT's performance, we developed the *Two-Panel Evaluation Methodology* for natural language generation and employed it in the most extensive and rigorous empirical evaluation ever conducted on an explanation system. In this study, KNIGHT constructed explanations on randomly chosen topics from the Biology Knowledge Base. A panel of domain experts was instructed to produce explanations on these same topics, and both KNIGHT's explanations and the explanations produced by this panel were submitted to a second panel of domain experts. The second panel then graded all of the explanations on several dimensions with an A–F scale. KNIGHT scored within approximately "half a grade" of the domain experts, and its performance exceeded that of one of the domain experts.

This paper is structured as follows. The task of explanation generation is characterized and the Biology Knowledge Base is described. A brief description of KNIGHT's knowledge-base access methods is followed by (1) a description of the EDP language, (2) KNIGHT's explanation planner, and (3) an overview of the realization techniques. The empirical evaluation is then discussed in some detail. The paper concludes with discussions of related work and future research directions.

2 The Task of Explanation Generation

Explanation generation is the task of extracting information from a formal representation of knowledge, imposing an organization on it, and realizing the information in text. An explanation system must be able to map from a formal representation of domain knowledge, i.e., one which can be used for automated reasoning, such as the predicate calculus, to a textual representation of domain knowledge. Because of the significant differences in formal and textual representational schemes, successfully bridging the gap between them is one of the major challenges faced by an explanation system.

To communicate complex ideas, an explanation system should be able to produce extended explanations such as those in Figure 1, which shows several explanations from the domain of botanical anatomy, physiology, and development. Note that each of these is a *multi-sentential* explanation; the first is a *multi-paragraph* explanation. These explanations are expository: in contrast to causal explanations produced by automated reasoning systems, they describe domain phenomena, such as anatomical structures and physiological processes.

Explanation generation is typically decomposed into two subtasks, *explanation planning* and *realization* (Figure 2). Explanation planning itself has two subtasks: *content determination*, in which knowledge structures are extracted from a knowledge base, and *organization*, in which the selected knowledge structures are arranged in a manner appropriate for communication in natural language. To perform these tasks successfully, an explanation planner must have access to *discourse knowledge*, which informs its decisions about the content and organization of textual explanations. The organizational aspect of discourse knowledge plays a particularly important role in the construction of extended explanations. Our work focuses on content determination and organization and deemphasizes issues in realization. However, to more cleanly evaluate the explanation planning work, we have developed a robust realization system that is built on FUF [15, 16], a unification-based implementation of a large systemic grammar.

2.1 Evaluation Criteria and Desiderata

Evaluating the performance of explanation systems is a critical and non-trivial problem. Although gauging the performance of explanation systems is inherently difficult, five evaluation criteria should be applied.

- *Coherence*: a global assessment of the overall quality of explanations generated by a system,
- *Content*: the extent to which the explanation's information is adequate and focused,
- *Organization*: the extent to which the information is well organized,
- *Writing style*: the quality of the prose, and
- *Correctness*: for scientific explanations, the extent to which the explanations are in accord with the established scientific record.

In addition to performing well on the evaluation criteria, if explanation systems are to make the difficult transition from research laboratories to fielded applications, they should exhibit two important properties, both of which significantly affect scalability. First, these systems' representation of

Question: What happens during *embryo sac formation*?

Response: Embryo sac formation is a kind of female gametophyte formation. During embryo sac formation, the embryo sac is formed from the megaspore mother cell. Embryo sac formation occurs in the ovule.

Embryo sac formation is a step of angiosperm sexual reproduction. It consists of megasporogenesis and embryo sac generation. During megasporogenesis, the megaspore mother cell divides in the nucellus to form 4 megaspores. During embryo sac generation, the embryo sac is generated from the megaspore.

Question: What happens during *pollen tube growth*?

Response: Pollen tube growth is a step of angiosperm sexual reproduction. It consists of sperm cell generation and sperm cell transport. During sperm cell generation, the pollen generative cell divides to form 2 angiosperm sperm cells. During sperm cell transport, 2 angiosperm sperm cells are transported from the pollen tube to the embryo sac.

Question: What is a *spore*?

Response: The spore is a kind of haploid cell. 4 spores are produced from the spore mother cell during sporogenesis. The spore divides to form 2 plant gametes during gametogenesis. Gametogenesis is a step of gametophyte development.

Question: What is a *root system*?

Response: The root system is part of the plant and is connected to the mainstem. It is below the hypocotyl and is surrounded by the rhizosphere. The subregions of the root system include the meristem, which is where root system growth occurs.

Figure 1: Explanations from the domain of biology produced by KNIGHT from the Biology Knowledge Base

- **Explanation Planning**: Constructing an *explanation plan*
  - **Content Determination**: Extracting structures from a representation of domain knowledge.
  - **Organization:** Arranging the selected knowledge structures.
- **Realization**: Translating an explanation plan to natural language.

Figure 2: The Task of Explanation Generation

discourse knowledge should be easily inspected and modified. To develop explanation systems for a broad range of domains, tasks, and question types, discourse-knowledge engineers must be able to create and efficiently debug the discourse knowledge that drives the systems' behavior. The second property that explanation systems should exhibit is robustness. Despite the complex and possibly mal-formed representational structures that an explanation system may encounter in its knowledge base, it should be able to cope with these structures and construct reasonably well-formed explanations.

2.2 Semantically Rich, Large-Scale Knowledge Bases

Given the state of the art in explanation generation, the field is now well positioned to explore what may pose its greatest challenge and at the same time may result in its highest payoff: generating explanations from semantically rich, large-scale knowledge bases. Large-scale knowledge bases encode information about domains that cannot be reduced to a small number of principles or axioms. For example, the field of human anatomy and physiology encompasses a body of knowledge so immense that many years of study are required to assimilate only one of its subfields, such as immunology. Large-scale knowledge bases are currently being constructed for many applications, and the ability to generate explanations from these knowledge bases for a broad range of tasks such as education, design and diagnosis is critical.

Large-scale knowledge bases whose representations are semantically rich are particularly intriguing. These knowledge bases consist of highly interconnected networks of (at least) tens of thousands of facts. Hence, they represent information not only about a large number of concepts but also about a large number of relationships that hold between the concepts. One such knowledge base is the **Biology Knowledge Base** [53], an immense structure encoding information about botanical anatomy, physiology, and development. One of the largest knowledge bases in existence, it is encoded in the KM frame-based knowledge representation language.[1] KM provides the basic functionalities of other frame-based representation languages and is accompanied by a graphical

[1] A detailed description of the semantics of the representation language may be found in Ch. 2 of [1].

user interface, KnEd, for entering, viewing, and editing frame-based structures [14].

The backbone of the Biology Knowledge Base is its taxonomy, which is a large hierarchical structure of biological objects and biological processes. In addition to the objects and processes, the taxonomy also includes the hierarchy of relations that may appear on concepts. The relation taxonomy provides a useful organizing structure for encoding information about "second order" relations, i.e., relations among all of the first order relations.

Figure 3 depicts a small portion of the Biology Knowledge Base's representation of *photosynthesis*. This is a typical fragment of its semantic network. Each of the nodes in this network is a concept, e.g. *photosynthesis*, which we frequently refer to as a "unit" or a "frame." Each of the arcs is a relation in the knowledge base. For example, the *transducer* for *photosynthesis* is the concept *chlorophyll*. We sometimes refer to these relations as "slots" or "attributes" and to the units that fill these slots, e.g., *chlorophyll*, as "values." In addition, we term a structure of the form ($Unit$ $Slot$ $Value$) as a "triple." The Biology Knowledge Base currently contains more than 180,000 explicitly represented triples, and its deductive closure is significantly larger.

We chose biology as a domain for three reasons. First, it required us to grapple with difficult representational problems. Unlike a domain such as introductory geometry, biology cannot be characterized by a small number of axioms. Second, biology is not a "single-task" subject. Unlike the knowledge bases of conventional expert systems, e.g., MYCIN [6], the Biology Knowledge Base is not committed to any particular task or problem-solving method. Rather, it encodes general knowledge that can support diverse tasks and methods such as tutoring students, performing diagnosis, and organizing reference materials. Finally, we chose biology because of the availability of domain experts at the University of Texas at Austin.

It is important to note that the authors and the domain experts entered into a "contractual agreement" with regard to representational structures in the Biology Knowledge Base. Under no circumstances could the designers of the explanation system request the knowledge engineers to alter the structures of the knowledge base other than for purposes of consistency or completeness. To eliminate all requests for representational modifications that would skew the knowledge base to the task of explanation generation, the authors entered into this agreement: they could request representational changes only if knowledge was inconsistent or missing. This facilitated a unique experiment in which the representational structures were not tailored for the task of explanation generation.

3 Accessing Semantically Rich, Large-Scale Knowledge Bases

To perform well, an explanation system must select from a knowledge base precisely that information needed to answer users' questions with coherent and complete explanations. Given the centrality of content determination for explanation generation, it is instructive to distinguish two types of content determination, both of which play key roles in an explanation system's behavior. "Local" content determination is the selection of relatively small knowledge structures, each of

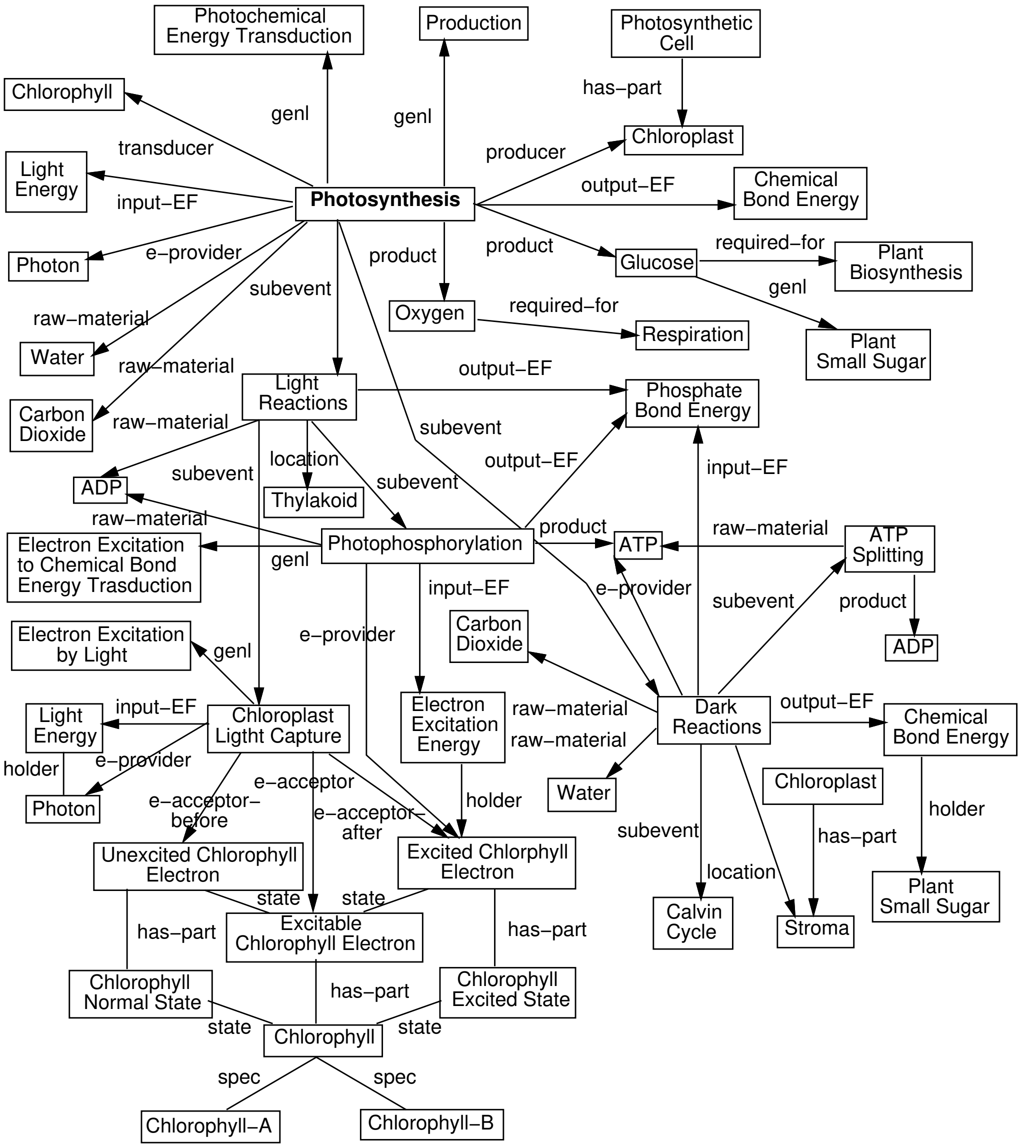

Figure 3: A sample representational structure from the Biology Knowledge Base

which will be used to generate one or two sentences; "global" content determination is the process of deciding which of these structures to include in an explanation.

By interposing a KB accessing system between an explanation planner, which performs global content determination, and a knowledge base, it is possible to keep an explanation planner at "arm's length" from the representation of domain knowledge. In addition, it can help build explanations that are coherent. Studies of coherence have focused on one aspect of coherence, *cohesion*, which is determined by the overall organization and realization of the explanation [13, 18, 19, 20, 26]. However, an equally important question is, "To insure coherence, how should the content of individual portions of an explanation be selected?" Halliday and Hassan [19] term this aspect of coherence *semantic unity*. There are at least two approaches to achieving semantic unity: either "packets" of propositions must be directly represented in the domain knowledge, or a KB accessing system must be able to extract them at runtime.

One kind of a semantic unit is a *view*. For example, the concept *photosynthesis* can be viewed as either a *production* process or an *energy transduction process*. Viewed as *production*, it would be described in terms of its *raw materials* and *products*: "During photosynthesis, a chloroplast uses water and carbon dioxide to make oxygen and glucose." Viewed as *energy transduction*, it would be described in terms of *input energy forms* and *output energy forms*: "During photosynthesis, a chloroplast converts light energy to chemical bond energy." Clearly, the view that is taken of a concept has a significant effect on the content that is selected for its description. If an explanation system could (a) invoke a knowledge base accessing system to select views, and (b) translate the views to natural language (Figure 4), it would be well on its way to producing coherent explanations.

As a building block for the KNIGHT explanation system, we designed and implemented a robust KB accessing system (Figure 5) that extracts *views* [5, 66, 57, 28, 50, 44, 64, 3, 1, 39, 59, 63] of concepts represented in a knowledge base. Each view is a coherent subgraph of the knowledge base describing the structure and function of objects, the change made to objects by processes, and the temporal attributes and temporal decompositions of processes. Each of the nine accessors in our library (Table 1) can be applied to a given concept—the "concept of interest"—to retrieve a view of that concept. There are three classes of Accessors: those that are applicable to all concepts (*As-Kind-Of* and *Functional*), those that are applicable to objects (*Partonomic-Connection* and *Sub-Structural*), and those that are applicable to processes (*Auxilliary-Process*—which includes *Causal*, *Modulatory*, *Temporal*, and *Locational* sub-types—*Participants*, *Core-Connection*, and *Sub-event*, and *Temporal-Step*).[2]

To illustrate, the *Participants* accessor extracts information about the "actors" of the given process. For example, some of the actors in the Photosynthesis process are Chloroplasts, Light, Chlorophyll, Carbon Dioxide, and Glucose (Figure 3). By specifying a *reference process*—the second

[2]In addition to the "top level" accessors, the library also provides a collection of some twenty "utility" accessors. These include procedures for extracting particular aspects of views previously constructed by the system. For example, the *Collect Actors* Accessor extracts the objects that serve as participants in a Participants view of process. For a more comprehensive description of the accessors, see [29].

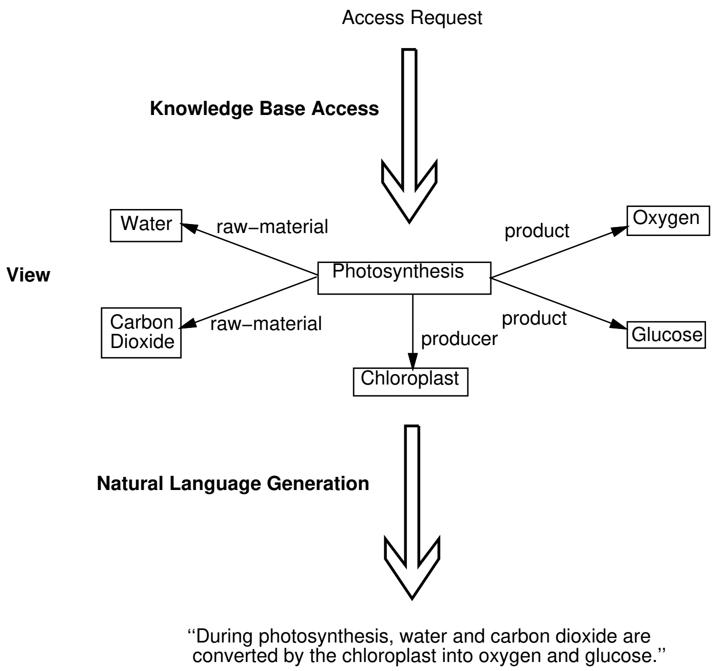

Figure 4: Accessing and translating a view of photosynthesis

argument to the *Participants* accessor—the external agent can request a view of the process from the perspective of the reference process. For example, if the system applies the *Participants* accessor with Photosynthesis as the concept of interest and Production as the reference process, then the accessor will extract information about the producer (Chloroplast), the raw materials (Water and Carbon Dioxide), and the products (Oxygen and Glucose). In contrast, if the system applies the *Participants* accessor with Photosynthesis as the concept of interest but with Energy Transduction as the reference process, then it would extract information about the transducer (Chlorophyll), the energy provider (a Photon), the input energy form (Light), and the output energy form (Chemical Bond Energy). By selecting different reference concepts, different information about a particular process will be returned.

In addition to coherence, robustness is also an important design criterion. The KB accessors achieve robust performance in four ways:

- *Omission Toleration:* They do not assume that essential information will actually appear on a given concept in the knowledge base.
- *Type Checking:* They employ a type checking system that exploits the knowledge base's taxonomy.
- *Error Handling:* When they detect an irregularity, they return appropriate error codes to the explanation planner.

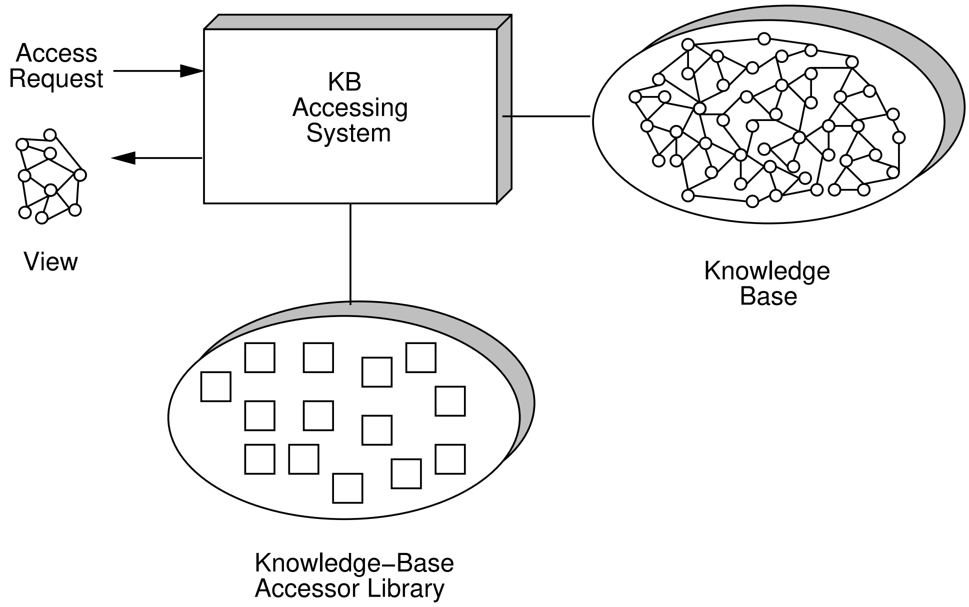

Figure 5: A KB accessing architecture

- *Term Accommodation:* They tolerate specialized (and possibly unanticipated) representational vocabulary by exploiting the relation taxonomy.

By using these techniques in tandem, we have developed a KB accessing system that has constructed several thousand views without failing. Moreover, the view types on which the accessors are based performed well in a preliminary empirical study [2], and evaluations of the KB accessors' ability to construct coherent views, as measured by domain experts' ratings of KNIGHT's explanations (Section 8), are encouraging.

4 A Programming Language for Discourse Knowledge

Since the time of Aristotle, a central tenet of rhetoric has been that a rich structure underlies text. This structure shapes a text's meaning and assists its readers in deciphering that meaning. For almost two decades, computational linguists have studied the problem of automatically inducing this structure from a given text. Research in explanation planning addresses the inverse problem: automatically constructing this structure by selecting facts from a knowledge base and subsequently using these facts to produce text. To automatically construct *explanation plans* (trees that encode the hierarchical structure of texts, as well as their content), an explanation system must possess knowledge about what characterizes a clear explanation. This *discourse knowledge* enables it to make decisions about what information to include in its explanations and how to organize the information.

It is important to emphasize the following distinction between discourse knowledge and ex-

KB Accessor	*Arguments*	*Description of View*
As-Kind-Of	*concept* *reference*	Finds view of *concept* as a kind of *reference* concept.
Auxiliary-Process	*process* *view-type*	Finds temporal, causal, or locational information about *process* as specified by *view-type*.
Participants	*process* *reference*	Finds "actor-oriented" view of *process* as viewed from the perspective of *reference* process.
Core-Connection	*process*	Finds the connection between *process* and a "core" process.
Functional	*object* *process*	Finds functional view of *object* with respect to *process*.
Partonomic-Connection	*object*	Finds the connection from *object* to a "superpart" of the object in the "partonomy."
Subevent	*process*	Finds view of "steps" of *process*.
Sub-Structural	*object*	Finds structural view of parts of *object*.
Temporal-Step	*process*	Finds view of *process* with respect to another process of which *process* is a "step."

Table 1: Library of Knowledge Base Accessors

planation plans. While discourse knowledge specifies the content and organization for a *class* of explanations, e.g., explanations of processes, explanation plans specify the content and organization for a *specific* explanation, e.g., an explanation of how photosynthesis produces sugar. Discourse-Knowledge engineers build representations of discourse knowledge, and this discourse knowledge is then used by a computational module to automatically construct explanation plans, which are then interpreted by a realization system to produce natural language.

The KB accessing system described above possesses discourse knowledge in the form of KB Accessors. Applying this discourse knowledge, the system retrieves views from the knowledge base. Although this ability to perform local content determination is essential, it is insufficient; given a query posed by a user, the generator must be able to choose multiple KB accessors, provide the appropriate arguments to these accessors, and organize the resulting views. Hence, in addition to discourse knowledge about local content determination, an explanation system that produces multi-paragraph explanations must also possess knowledge about how to perform global content determination and organization. This section sets forth two design requirements for a representation of discourse knowledge, describes the *Explanation Design Package* (EDP) formalism, which was designed to satisfy these requirements, and discusses how EDPs can be used to encode discourse knowledge.

4.1 Requirements for a Discourse Knowledge Representation

A representation of discourse knowledge should satisfy two requirements: It should be expressive, and it should facilitate efficient representation of discourse knowledge by discourse-knowledge engineers. Each of these considerations are discussed in turn, followed by a representation that satisfies these criteria.

Expressiveness. A representation of discourse knowledge must permit discourse-knowledge engineers to state how an explanation planner should

- select propositions from a knowledge base,
- control the amount of detail in an explanation, i.e., if a user requests that terse explanations be generated, the explanation planner should select only the most important propositions
- consider contextual conditions when determining which propositions to include,
- order the propositions, and
- group the propositions into appropriate segments, e.g., paragraphs.

The first three aspects of expressiveness are concerned with content determination. To effectively express what content should be included in explanations, a representation of discourse knowledge should enable discourse-knowledge engineers to encode specifications about how to choose propositions about particular topics, the importance of those topics, and under what conditions the propositions associated with the topics should be included. These "inclusion conditions" govern the circumstances under which the explanation planner will select particular classes of propositions from the knowledge base when constructing an explanation.

For example, a discourse-knowledge engineer might express the rule: "The system should communicate the location of a process if and only if the user of the system is familiar with the object where the process occurs." As the explanation planner uses this knowledge to construct a response, it can determine if the antecedent of the rule ("the user of the system is familiar with the object where the process occurs") is satisfied by the current context; if the antecedent is satisfied, then the explanation planner can include in the explanation the subtopics associated with the rule's consequent.

The final two aspects of expressiveness (ordering and grouping of propositions) are concerned with organization. To encode organizational knowledge, a representation of discourse knowledge should permit discourse-knowledge engineers to encode topic/subtopic relationships. For example, the subtopics of a process description might include (1) a categorical description of the process (describing taxonomically what kind of process it is), (2) how the "actors" of the process interact, and (3) the location of the process.

A representation should be sufficiently expressive that it can be used to encode the kinds of discourse knowledge discussed above, and it should be applicable to representing discourse knowledge for a broad range of discourse genres and domains. However, discourse knowledge does not specify what syntactic structure to impose on a sentence; nor does it lend any assistance in making decisions about matters such as pronominalization, ellipsis, or lexical choice. These decisions are delegated to the realization system.

Discourse-Knowledge Engineering. For a given query type, domain, and task, a discourse-knowledge engineer must be able to represent the discourse knowledge needed by an explanation system for responding to questions of that type in that domain about that task. Pragmatically, to represent discourse knowledge for a broad range of queries, domains, and tasks, a formalism must facilitate *efficient* representation of discourse knowledge. Therefore, important goals for the design of a discourse formalism are ease of reuse and ease of modification.

For example, to build an explanation system for the domain of physics, a discourse-knowledge engineer could either build an explanation system *de novo* or modify an existing system. On the face of it, the second alternative involves less work and is preferable, but designing explanation systems that can be easily modified is a non-trivial task. In the case of physics, a discourse-knowledge engineer may need to modify an existing explanation system so that it can produce explanations that are appropriate for mathematical explanations. To do so, the discourse-knowledge engineer would ideally take an off-the-shelf explanation generator and add discourse knowledge about how to explain mathematical interpretations of the behavior of physical systems. Because of the central role played by discourse-knowledge engineers, a representation of discourse knowledge should be designed to minimize the effort required to understand, modify, and represent new discourse knowledge.

4.2 Explanation Design Packages

Explanation Design Packages emerged from an effort to accelerate the representation of discourse knowledge without sacrificing expressiveness. Our previous explanation generators employed a representation of discourse knowledge that consisted of collections of operators, coded directly in Lisp, for determining the content and organization of explanations [30, 32, 33, 31]. Although this approach worked well for small prototypes, it proved unsatisfactory for building fully functioning explanation systems. In particular, it was very difficult to maintain and extend discourse knowledge expressed directly in code.

EDPs language give discourse-knowledge engineers the proper abstractions for specifying the content and organization of explanations. They combine a frame-based representation language with embedded procedural constructs. To mirror the structure of expository texts, an EDP contains a hierarchy of nodes, which provides the "global organization" of explanations. EDPs are schema-like [41, 52] structures that include constructs found in traditional programming languages. Just as prototypical programming languages offer conditionals, iterative control structures, and procedural

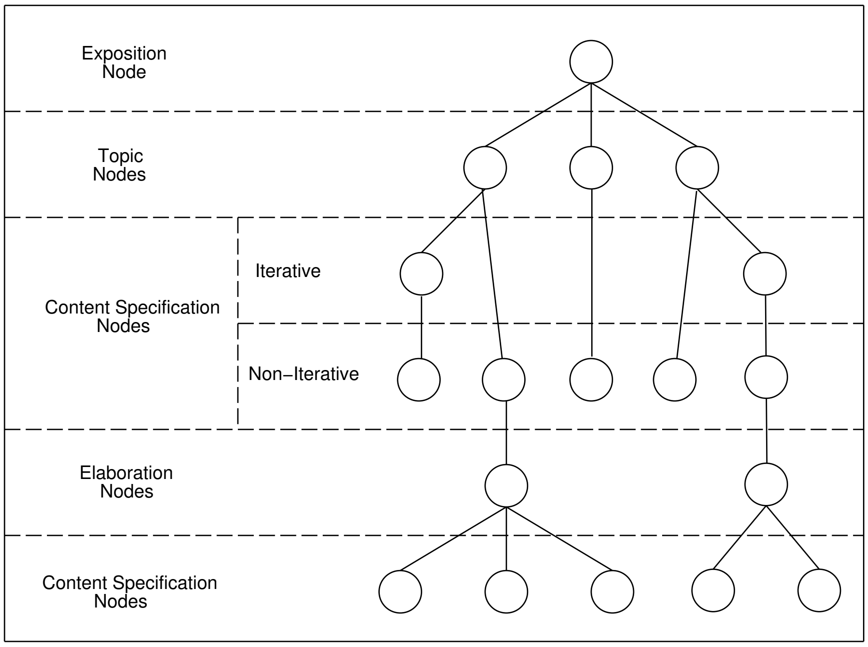

Figure 6: The Layers of Explanation Design Packages

abstraction, EDPs offer discourse-knowledge engineers counterparts of these constructs that are precisely customized for explanation planning. Moreover, each EDP names multiple KB accessors, which are invoked at explanation planning time.

Because EDPs are frame-based, they can be easily viewed and edited by knowledge engineers using the graphical tools commonly associated with frame-based languages. The EDP formalism has been implemented in the KM frame-based knowledge representation language, which is the same representational language used in the Biology Knowledge Base. Because KM is accompanied by a graphical user interface, discourse-knowledge engineers are provided with a development environment that facilitates EDP construction. This has proven to be very useful for addressing a critical problem in scaling up explanation generation: maintaining a knowledge base of discourse knowledge that can be easily constructed, viewed, and navigated by discourse-knowledge engineers.

EDPs have several types of nodes (Figure 6), where each type provides a particular set of attributes to the discourse-knowledge engineer (Table 2). Note that content specification nodes may have elaboration nodes as their children, which in turn may have their own content specification nodes. This "recursive" appearance of content specification nodes permits a discourse-knowledge engineer to construct arbitrarily deep trees. In general, a node of a particular type in an EDP is used by the explanation planner to construct a corresponding node in an explanation plan. We discuss the salient aspects of each type of node below.[3]

[3]Representational details of EDPs are discussed in [29].

Node Type	*Attributes*	*Attribute Value(s)*
Exposition	Children	⟨Topics⟩
Topic	Children	⟨Content Specifications⟩
	Centrality	{Low, Medium, High}
	Inclusion Condition	⟨Variable Boolean Expression⟩
	Local Variables	(⟨Var⟩ , ⟨Variable Expr.⟩) Pairs
Content Specification	Children	{⟨Content Spec's⟩, ⟨Elaborations⟩}
	Content Specification Template	⟨Variable Expression with KB Accessor⟩
	Iteration Type	{Non-Iter., Iter., Conditional-Iter.}
	Iterate-Over Template	⟨Variable Expression⟩
	Loop Variable	⟨Var⟩
	Iteration Condition	⟨Variable Boolean Expression⟩
	Local Variables	(⟨Var⟩ , ⟨Variable Expr.⟩) Pairs
Elaboration	Children	⟨Content Specifications⟩
	Centrality	{Low, Medium, High}
	Inclusion Condition	⟨Variable Boolean Expression⟩
	Local Variables	(⟨Var⟩ , ⟨Variable Expr.⟩) Pairs

Table 2: EDP Node Attributes

Exposition Nodes. An exposition node is the top-level unit in the hierarchical structure and constitutes the highest-level grouping of content. For example, the exposition node of the *Explain-Process* EDP has four children, *Process Overview*, *Output-Actor-Fates*, *Temporal Info*, and *Process Details*, each of which is a topic node. Both the order and grouping of the topic nodes named in an exposition node are significant. The order specifies the linear left-to-right organization of the topics, and the grouping specifies the paragraph boundaries. The content associated with topic nodes that are grouped together will appear in a single paragraph in an explanation.

Topic Nodes. Topic nodes are "subtopics" of exposition nodes. Each topic node includes a representation of the conditions under which its content should be added to an explanation. A topic node has the *atomic inclusion* property:

> An explanation planner can make an "atomic" decision about whether to include—or exclude—all of the content associated with a topic node, i.e., either *all* of the content associated with a topic node will be included in an explanation or *none* of it will.

To enable the explanation planner to make fine-grained decisions about content inclusion, discourse-knowledge engineers can group closely related content together under the same topic node. The degenerate case is one in which all content is placed under a single topic node; this will force the

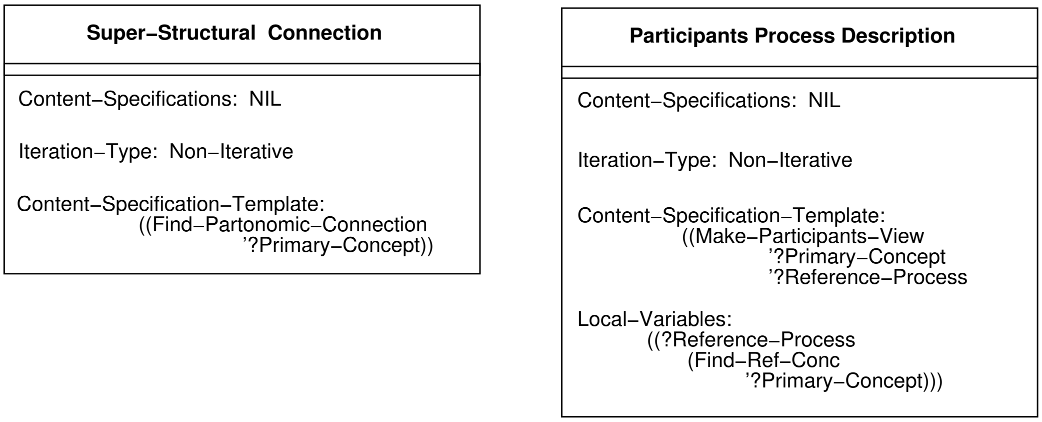

Figure 7: Example content specification nodes

discourse planner to include either all of the content or none of it.

An important aspect of discourse knowledge is the relative importance of subtopics with respect to one another. If an explanation's length must be limited—such as when a user has employed the verbosity preference parameter to request terse explanations—an explanation planner should be able to decide at runtime which propositions to include. EDPs permit discourse-knowledge engineers to specify the relative importance of each topic by assigning a qualitative value (`Low`, `Medium`, or `High`) to its centrality attribute.

Another important aspect of representing discourse knowledge is the ability to encode the conditions under which a group of propositions should be included in an explanation. Discourse-knowledge engineers can express these *inclusion conditions* as predicates on the knowledge base and on a user model (if one is employed). For example, he or she should be able to express the condition that the content associated with the *Output Actor Fates Topic* should be included only if the process being discussed is a conversion process. Inclusion conditions are expressed as boolean expressions that may contain both built-in user modeling predicates and user-defined functions.

Content Specification Nodes. Content specification nodes house the high-level specifications for extracting content from the knowledge base. To fulfill this function, they provide constructs known as *content specification expressions*. These expressions are instantiated at runtime by the explanation planner, which then dispatches the Knowledge Base Accessors named in the expressions to extract propositions from the knowledge base. Content specification expressions reside in content specification nodes, e.g., Figure 7. When creating content specification expressions, the discourse-knowledge engineer may name any Knowledge Base Accessor in the KB Accessor Library. For example, the *Super Structural Connection* content specification in Figure 7 names a KB Accessor called *Find Partonomic Connection*, and the *Process Participants Description* content specification names the *Make-Participants-View* Accessor.

Although the discourse-knowledge engineer may write arbitrarily complex specification expressions in which function invocations are deeply nested, these expressions can become difficult to understand, debug, and maintain. Just as other programming languages provide local variables, e.g., the binding list of a `let` statement in Lisp, so do content specification nodes. Each time a discourse-knowledge engineer creates a local variable, he or she creates an expression for computing the value of the local variable at runtime. For example, the *Process Participants Description* content specification in Figure 7 employs a local variable *?Reference Process*. The content specification expression associated with *?Reference Process* names the KB Accessor *Find Ref Conc* and the global variable *?Primary-Concept*. Local variables provide a means for decomposing more complex content specification expressions into simpler ones.

There are three types of content specification nodes: *Non-Iterative*, *Iterative*, and *Conditional Iterative*. The discourse-knowledge engineer indicates the type of a content specification node in its iteration-type attribute. Thus far we have been discussing Non-Iterative nodes, but some discourse knowledge is inherently iterative. For example, suppose a discourse-knowledge engineer is representing discourse knowledge to be used by an explanation planner for generating explanations about objects. The discourse-knowledge engineer may need to represent the specification that, for a given object O,

for each part P of O
    for each function F of P
        describe how F is a function of P

This specification involves a doubly nested iteration: the outer statement iterates over the parts of an object; the inner statement iterates over the functions of one of the object's parts. In general, an important aspect of representing discourse knowledge is the ability to specify that the discourse planner should repeat an operation for each entity in a set.[4] By including constructs for both standard and conditional iteration, EDPs provide discourse-knowledge engineers with a formalism that enables them to represent commonly occurring iterative discourse patterns.

Elaboration Nodes. Elaboration nodes specify optional content that may be included in explanations. They are structurally identical to topic nodes, i.e., they have exactly the same attributes, and the children of elaboration nodes are content specifications. Elaboration nodes and topic nodes differ only in function; topic nodes are used to specify the primary content of explanations, and elaboration nodes are used to specify supplementary content.

4.3 Developing Task-Specific EDPs

A discourse-knowledge engineer can use EDPs to encode discourse knowledge for his or her application. In our work, we focused on two types of texts that occur in many domains: *process descriptions*

[4]Note that the set can only be specified abstractly by the discourse-knowledge engineer, e.g., "the functions of some structure." At runtime, the explanation planner will materialize the set of entities and iterate over them.

and *object descriptions*. For example, in biology, one encounters many process-oriented descriptions of physiological and reproductive mechanisms, as well as many object-oriented descriptions of anatomy. In the course of our research, we studied many passages in biology textbooks. These passages focused on explanations of the anatomy, physiology, and reproduction of plants. Some explanations were very terse, e.g., those that occurred in glossaries, whereas some were more verbose, e.g., multi-page explanations of physiological processes. Most of the texts also contained information about other aspects of botany, e.g., experimental methods and historical developments; these were omitted from the analysis.

We manually "parsed" each passage into a discourse tree. The discourse trees were expressed in an informal language centering around viewpoints [59, 56]. The viewpoints were in turn expressed in an informal language of structure, function, and process which is commonly found in the discourse literature, e.g., [41, 35, 52]. Our final step was to generalize the most commonly occurring patterns into abstractions that covered as many aspects of the passages as possible. After generalizing the commonly occurring patterns into abstractions, we encoded the abstractions in two Explanation Design Packages. These EDPs can be used by an explanation planner to generate explanations about the processes and objects of physical systems. Although we will not discuss the details of these EDPs here, it is instructive to examine their structure and function.[5]

The *Explain-Process* EDP (Figure 8) can be used by the explanation planner to generate explanations about the processes that physical objects engage in. For example, given a query about how a biological process such as Embryo Sac Formation is carried out, the explanation planner can apply the *Explain-Process* EDP to construct an explanation plan that houses the content and organization of the explanation. The *Explain-Process* EDP has five primary topics

- *Process Significance*: Explains the significance of a process with respect to the role it plays in a "parent" process;
- *Process Overview*: Explains how a process fits into a taxonomy, discusses the role played by its actors, and discusses where it occurs;
- *Process Details*: Explains the steps of a process;
- *Temporal Attributes*: Explains how a process is related temporally to other processes;
- *Output Actor Fates*: Discusses how the "products" of a process are used by other processes.

As computational linguists have known for many years, formally characterizing texts is a very difficult, time-consuming, and error-prone process. Because any initial discourse representation effort by necessity must be considered only a beginning, the next step was to incrementally revise the EDPs. The EDPs were used to automatically construct hundreds of explanations: the explanation planner used the EDPs to construct explanation plans, and the realization system translated these plans to natural language. The resulting explanations were presented to our domain expert, who

[5] All of the example EDP nodes in this section were taken from these EDPs.

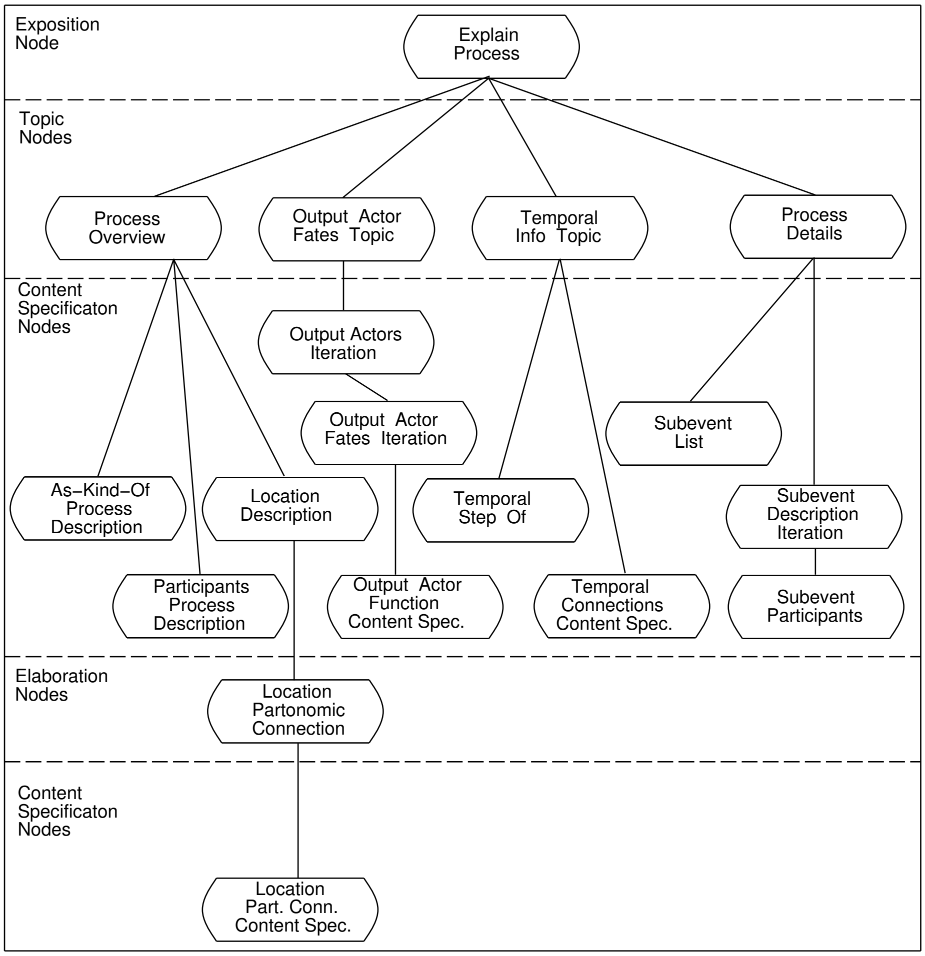

Figure 8: The Final Version of the Explain-Process Explanation Design

critiqued both their content and organization, and we used these critiques to incrementally revise the EDPs. After many passes through the critiquing and revision phases, EDPs were devised that produced clear explanations meeting with the domain expert's approval.

5 Planning Explanations

Explanation planning is the task of determining the content and organization of explanations. We have designed an architecture for explanation generation and implemented a full-scale explanation generator, KNIGHT,[6] that is based upon this architecture.

[6]All of the explanation planning algorithms, as well as the KB Accessors, were implemented in Lucid Common Lisp on a DEC Station 5000.

5.1 An Architecture for Explanation Generation

Explanation generation begins when the user poses a query, which includes a verbosity specification that comes in the form of a qualitative rating expressing the desired length of the explanation (Figure 9). The query interpreter—whose capabilities have been addressed only minimally in our work—translates the query to a canonical form, which is passed, along with the verbosity specification, to the explanation planner.

The heart of an explanation generator is its explanation planner. Explanation planning is a synthetic task in which multiple resources are consulted to assemble data structures that specify the content and organization of explanations. KNIGHT's planner uses the following resources:

- Domain knowledge representation: Biology Knowledge Base;
- Discourse knowledge representation: Explanation Design Packages;
- Domain knowledge access methods: KB accessing system;
- Context representation: Overlay user model.[7]

The explanation planner invokes the EDP Selector, which chooses an Explanation Design Package from the EDP Library. The explanation planner then applies the EDP by conducting an in-order traversal of its hierarchical structure. For each node in the EDP, the planner determines if it should construct a counterpart node in the explanation plan it is building. (Recall that the topic nodes and elaboration nodes of an EDP are instantiated only when their conditions are satisfied.) In this manner, the planner assembles an explanation plan that is isomorphic to the EDP. As the plan is constructed, the explanation planner updates the user model to reflect the contextual changes that will result from explaining the views in the explanation plan, attends to the verbosity specification, and invokes KB Accessors to extract information from the knowledge base. Recall that the Accessors return "views," which are subgraphs of the knowledge base. The planner attaches the views to the explanation plan; they become the plan's leaves. Planning is complete when the explanation planner has traversed the entire EDP.

The planner passes the resulting explanation plan to the realization component (Section 6) for translation to natural language. The views in the explanation plan are grouped into *paragraph clusters*. After some "semantic polishing" to improve the content for linguistic purposes, the realization component translates the views in the explanation plan to sentences. The realization system collects into a paragraph all of the sentences produced by the views in a particular paragraph cluster. Explanation generation terminates when the realization component has translated all of the views in the explanation plan to natural language.

[7] As the planner constructs explanation plans, it consults the user model. User-sensitive explanation generation is not addressed in this paper. For a discussion of this work, see [33, 29].

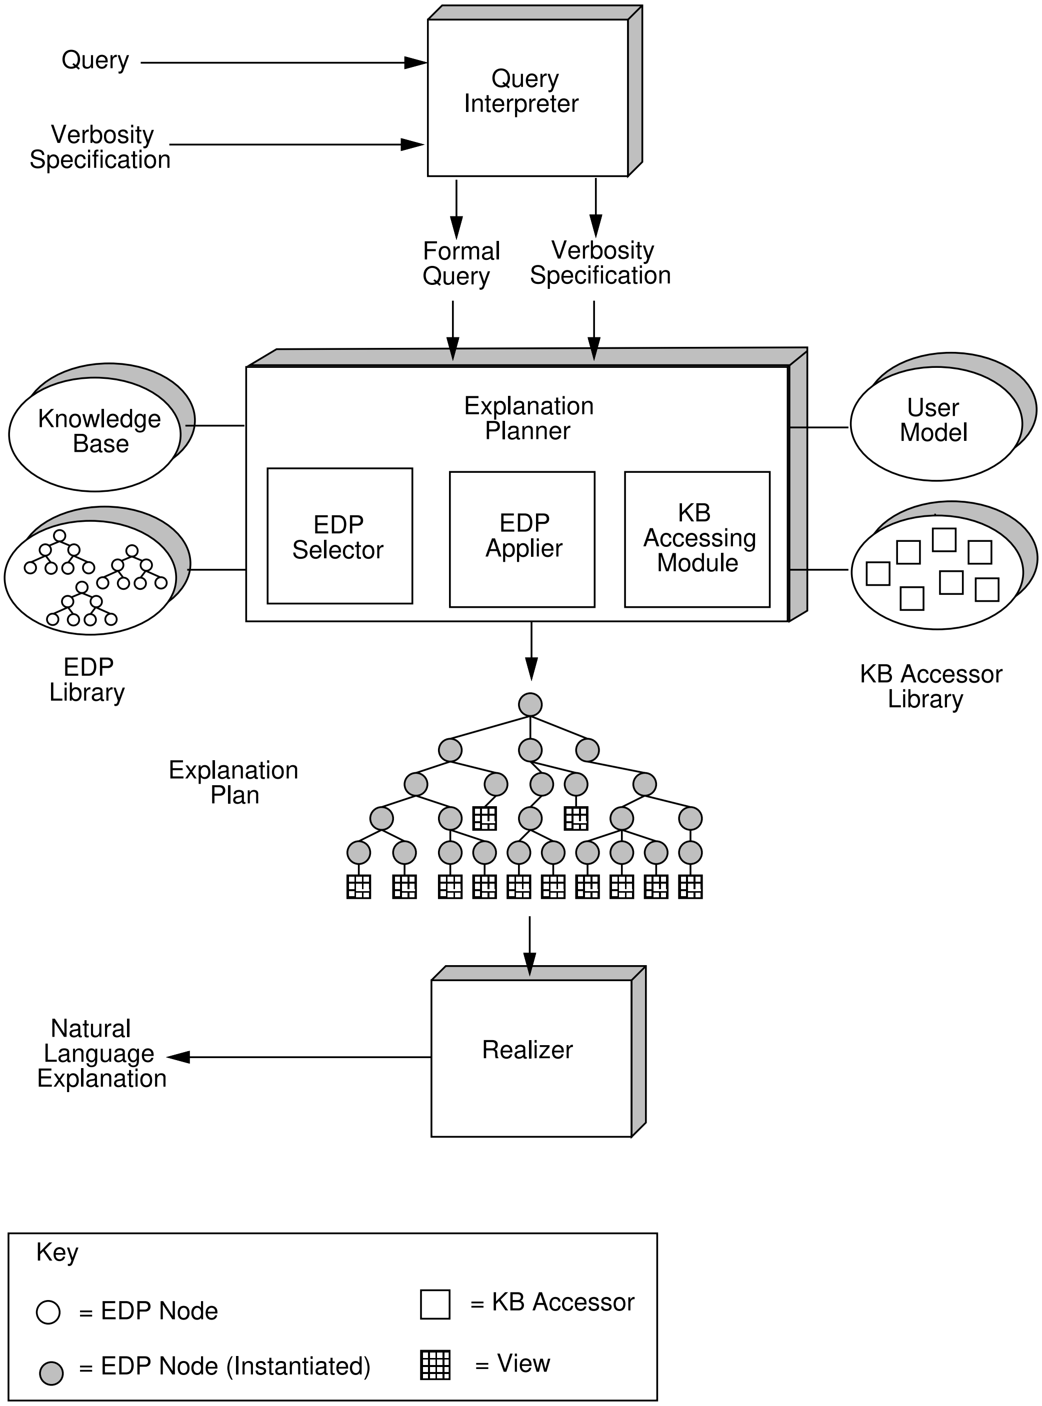

Figure 9: An Architecture for Explanation Generation

EXPLAIN (*Query-Type*, *Concept*, *Verbosity*)

if legal-query (*Query-Type*, *Concept*, *Verbosity*) **then**
  EDP ← select-edp (*Query-Type*)
  EDP-Exposition-Node ← get-root (*EDP*)
  New-Exposition-Node ← construct-node (*EDP-Exposition-Node*)
  Explanation-Plan ← apply-edp (*EDP-Exposition-Node*,
                                  New-Exposition-Node, *Verbosity*, **nil**)
  Explanation-Leaves ← linearize (*Explanation-Plan*)
  realize (*Explanation-Leaves*)

Figure 10: The EXPLAIN Algorithm

5.2 The Explanation Planning Algorithms

The EXPLAIN algorithm (Figure 10) is supplied with a query type (e.g., `Describe-Process`), a primary concept (e.g., `Embryo-Sac-Formation`), and a verbosity specification (e.g., `High`). Its first step is to select an appropriate EDP. Recall that the EDP Library has an indexing structure that maps a query type to the EDP that can be used to generate explanations for queries of that type. This indexing structure permits EDP selection to be reduced to a simple look-up operation. For example, given the query type *Describe Process*, the EDP Selector will return the *Explain-Process* Explanation Design Package. The planner is now in a position to apply the selected EDP to the knowledge base. The APPLY EDP algorithm takes four arguments: the exposition node of the EDP that will be applied, a newly created exposition node which will become the root of the explanation plan that will be constructed, the verbosity specification, and the loop variable bindings.[8]

The planner first locates the root of the selected EDP, which is an exposition node. Next, it creates the corresponding exposition node for the soon-to-be-constructed explanation plan. It then invokes the APPLY EDP algorithm, which is given the exposition node of the EDP to be applied, the newly created exposition node that will become the root of the explanation plan, the verbosity, and a list of the loop variable bindings.[9] The APPLY EDP algorithm (Figure 11) and the algorithms it invokes conduct an in-order traversal of the hierarchical structure of the EDP to build an explanation plan. Its first action is to obtain the children of the EDP's exposition node; these are the topic nodes of the EDP. For each topic node, the EDP Applier constructs a new (corresponding) topic node for the evolving explanation plan. The Applier must then weigh several factors in its decision about whether to include the topic in the explanation

- *Inclusion*: the inclusion condition associated with the topic;
- *Centrality*: the centrality rating that the discourse-knowledge engineer has assigned to the

[8] APPLY EDP is a recursive algorithm. For top-level invocations, this latter parameter will always be `nil`.
[9] The loop variable bindings are used for the EDPs' iteration construct.

```
APPLY-EDP (EDP-Exposit-Node, New-Exposit-Node, Verbosity,
           Loop-Var-Bindings)

Children-of-EDP-Exposition-Node
       ← get-children (EDP-Exposition-Node)
for each EDP-Topic-Node in EDP-Exposition-Node-Children do
   New-Topic-Node ← construct-node (EDP-Topic-Node)
   Inclusion-Condition-Expression ← get-condition (EDP-Topic-Node)
   Instantiated-Inclusion-Condition
             ← instantiate (Inclusion-Condition-Expression,
                            EDP-Topic-Node,
                            New-Topic-Node)
   Inclusion-Condition-Evaluation
           ← eval (Instantiated-Inclusion-Condition)
   Centrality ← get-centrality (EDP-Topic-Node)
   Include-Topic? ← compute-inclusion (Inclusion-Condition-Evaluation,
                                       Centrality,
                                       Verbosity)
   if Include-Topic? then
     Children-of-EDP-Topic-Node ← get-children (EDP-Topic-Node)
     for each EDP-Content-Specification-Node
         in Children-of-EDP-Topic-Node do
         determine-content (EDP-Content-Specification-Node,
                            New-Topic-Node,
                            Verbosity,
                            Loop-Var-Bindings)
```

Figure 11: The EDP Application Algorithm

topic;

- *Verbosity*: the verbosity specification supplied by the user.

If the inclusion condition evaluates to FALSE, the topic should be excluded regardless of the other two factors. Otherwise, the COMPUTE INCLUSION algorithm must consider the topic's importance and the amount of detail requested and will include the topic in the following circumstances: the verbosity is `High`; the verbosity is `Low` but the topic's centrality has been rated as `High` by the discourse knowledge engineer; or the verbosity is `Medium` and the topic's centrality has been rated as `Medium` or `High`.

When the COMPUTE INCLUSION algorithm returns TRUE, the Applier obtains the children of the EDP's topic. These are its content specification nodes. For each of the topic's content specification nodes, the Applier invokes the DETERMINE CONTENT algorithm, which itself invokes

KB Accessors named in the EDP's content specification nodes. This action extracts views from the knowledge base and attaches them to the explanation plan.

Rather than merely returning a flat list of views, the EXPLAIN algorithm examines the paragraph specifications in the nodes of the EDP it applied. The paragraph specifications of a given node organize the children of that node into *paragraph clusters*. The order of the paragraph clusters controls the global structure of the final textual explanation; the order of the views in each paragraph cluster determines the order of sentences in the final text.[10] Finally, the EXPLAIN algorithm passes the paragraph clusters to the REALIZE algorithm, which translates them to natural language.

5.3 The Products of Explanation Planning

Planning is complete when the APPLY EDP algorithm has traversed the entire EDP and constructed an explanation plan. To assist the discourse-knowledge engineer, explanation plan nodes are self-documenting. Each design decision is recorded directly in the attributes of the nodes themselves. When an explanation plan node is created, the attributes of the corresponding node in the Explanation Design Package are copied over to the new node in the explanation plan. In addition, the explanation planner records in explanation plans all variable bindings, instantiated content specification expressions, instantiated inclusion condition expressions, instantiated iteration condition expressions, flags indicating that KB Accessors have detected inappropriate structures in the knowledge base, and error annotations produced by KB Accessors indicating the nature of inappropriate knowledge-base structures. By providing this information in a structured design history, explanation plans significantly assist discourse-knowledge engineers: it permits them to rapidly pinpoint problems in the EDPs which were used to automatically construct the explanation plans.

The explanation plan formalism is represented in KM, the same frame-based representation language that is used for the Biology Knowledge Base and for EDP nodes. Each node in explanation plan is a dynamically created frame. Because KM is accompanied by a graphical user interface for viewing frame-based structures, discourse-knowledge engineers can easily view and navigate explanation plans. Hence, when discourse-knowledge engineers create new EDPs or modify existing ones, they can use the graphical user interface to inspect efficiently the results of explanation planning. They can then quickly locate deficiencies in explanation plans and use this information to revise the EDPs.

6 Realization

The explanation planner should be viewed as an automatic specification writer: its task is to write specifications for the realization component, which interprets the specifications to produce natural language. Although our work focuses on the design, construction, and evaluation of explanation

[10] The realization algorithm treats these groupings as suggestions which may be overridden in extenuating circumstances.

planners, by constructing a full-scale natural language generator, it becomes possible to conduct a "pure" empirical evaluation of explanation planners. Without a realization component, the plans produced by an explanation planner would need to be manually translated to natural language, which would raise many questions about the purity of the experiments. We therefore designed and implemented a full-scale realization component.[11]

Realization can be decomposed into two subtasks: *functional realization*, constructing functional descriptions from message specifications supplied by a planner; and *surface generation*, translating functional descriptions to text. *Functional descriptions* encode both semantic information (case assignments) and structural information (phrasal constituent embeddings). Syntactically, a functional description is a set of attribute and value pairs `(a v)` (collectively called a *feature set*), where `a` is an attribute (a feature) and `v` is either an atomic value or a nested feature set.[12] To illustrate, Figure 12 depicts a sample functional description. The first line, `(cat clause)`, indicates that what follows will be some type of verbal phrase, in this case a sentence. The second line contains the keyword `proc`, which denotes that everything in its scope will describe the structure of the entire verbal phrase. The next structure comes under the heading `partic`; this is where the thematic roles of the clause are specified. In this instance, one thematic role exists in the main sentence, the `agent` (or subject), which is further defined by its lexical entry and a modifying prepositional phrase indicated by the keyword `qualifier`.

The structure beginning with `circum` creates the subordinate infinitival purpose clause. It has two thematic roles, subject and object. The subject has a pointer to identify itself with the subject of the main clause while the object contains a typical noun phrase. The feature set for the `circum` clause indicates the wide range of possibilities for placement of the clause as well as for introducing additional phrasal substructures into the purpose clause.

KNIGHT employs the FUF surface realization system [15, 16] to translate functional descriptions into text. Developed by Elhadad and his colleagues at Columbia, FUF is accompanied by an extensive, portable English grammar, which is "the result of five years of intensive experimentation in grammar writing" and is currently the largest "generation grammar" in existence.

To construct functional descriptions from views extracted from a knowledge base, KNIGHT employs a functional realization system [7]. Given a view, the functional realizer uses its knowledge of case mappings, syntax, and lexical information to construct a functional description, which it then passes to the FUF surface generator. The functional realizer consists of five principal components:

- *Lexicon:* Physically distributed throughout the knowledge base; each concept frame has access to all of the lexical information relevant to its own realization.

[11] During the past few years, we have developed a series of realization systems. The first realizer, which was designed and implemented by the first author, was a template-based generator. The second realizer, which was designed by Kathy Mitchell and the authors, used the Penman surface generator. Her implementation of the second system is described in [45]. The third realizer is described briefly in this section; it was designed by the first author and C. Callaway.

[12] Functional descriptions may also employ syntactic sugar for purposes of legibility.

```
((cat clause)
 (proc ((type material) (lex ``reproduce'')))
 (partic ((agent ((cat common)
                  (lex ``spore'')
                  (qualifier ((cat pp)
                              (prep === ``from'')
                              (np ((cat common)
                                   (lex ``cell'')
                                   (classifier ((cat noun-compound)
                                                (classifier === ``megaspore'')
                                                (head === ``mother'')))
                                   (qualifier ((cat pp)
                                               (prep === ``in'')
                                               (np ((cat common)
                                                    (lex ``sporangium'')
                                                   ))))))))))))
 (circum ((purpose ((cat clause) (position end)
                    (keep-for no) (keep-in-order no)
                    (proc ((type material) (lex ``form'')))
                    (partic ((agent ((semantics {partic agent semantics})))
                             (affected ((cat common)
                                        (lex ``gamete'') (definite no)
                                        (classifier === ``plant'')
                                        (describer === ``haploid'')
                                        (cardinal ((value 4) (digit no))))
                                       )))))
          (time ((time-type ``during'')
                 (cat common)
                 (describer === ``male'')
                 (classifier === ``gametophyte'')
                 (lex ``generation''))))))
```

Figure 12: A Functional Description

- *Functional Description Skeleton Library:* Contains a large number of FD-Skeletons, each of which encodes the associated syntactic, semantic, and role assignments for interpreting a specific type of message specification.
- *Functional Description Skeleton Retriever:* Charged with the task of selecting the correct Functional Description Skeleton from the skeleton library.
- *Noun Phrase Generator:* Responsible for drawing lexical information from the Lexicon to create a self-contained functional description representing each noun phrase required by the FD-Skeleton processor
- *Functional Description Skeleton Processor:* Gathers all of the available information from the FD-Skeleton, the lexicon, and the noun phrase generator; produces the final functional description.

When the functional realizer is given a view, its first task is to determine the appropriate FD-Skeleton to use. Once this is accomplished, the FD-Skeleton is passed along with the message specification to the FD-Skeleton processor. The FD-Skeleton processor first determines if each of the essential descriptors are present; if any of these tests fail, it will note the deficiency and abort. If the message is well-formed, the FD-Skeleton processor passes each realizable concept unit found on the message specification to the noun phrase generator, which uses the lexicon to create a functional description representing each concept unit. The noun phrase generator then returns each functional description to the FD-Skeleton processor, which assigns case roles to the (sub)-functional descriptions. The resulting functional description, which encodes the functional structure for the entire content of the message specification, is then passed to the surface realizer, which constructs the final text.[13]

7 Example Behavior

To illustrate the behavior of the system, consider the concept of *embryo sac formation.* Figure 13 depicts the semantic network in the Biology Knowledge Base that represents information about *embryo sac formation.* When KNIGHT is given the task of explaining this concept,[14] it applies the *Explain-Process* EDP (Figure 8).

KNIGHT first finds the topics of the *Explain Process* exposition node, which are *Process Overview*, *Output Actor Fates*, *Temporal Information*, and *Process Details.* During its traversal of this tree, it begins with *Process Overview*, which has a `High` centrality rating and an inclusion condition of TRUE. KNIGHT executes the COMPUTE INCLUSION algorithm with the given verbosity of `High`, which returns TRUE, i.e., the information associated with the topic should be included.

Hence, it now begins to traverse the children of this topic node, which are the *As Kind Of Process Description*, *Process Participants*, and *Location Description* content specification nodes. For the *As Kind Of Process Description*, it computes a value for the local variable *?Reference Concept*, which returns the value *female gametophyte formation.* It then instantiates the content specification template on *As Kind Of Process Description*, which it then evaluates. This results in a call to the *As-Kind-Of* KB Accessor, which produces a view. Similarly, KNIGHT instantiates the content specification expressions of *Process Participants Description* and *Location Description*, which also cause KB Accessors to be invoked; these also return views. Next KNIGHT visits the *Location Partonomic Connection* node, which is an elaboration of *Location Description.* However, because its inclusion condition is not satisfied, this branch of the traversal halts.

Next, KNIGHT visits each of the other topics of the *Explain Process* exposition node: *Temporal Information* and *Process Details.* Because it was given a `High` verbosity specification, both of these are used to determine additional content. The resulting explanation plan is shown in Figure 14. Note that the *Subevent Description Iteration*, together with the *Subevent Participants* nested con-

[13] More comprehensive descriptions of KNIGHT's realization component may be found in [29] and [7].

[14] In this example, KNIGHT is given a HIGH verbosity specification.

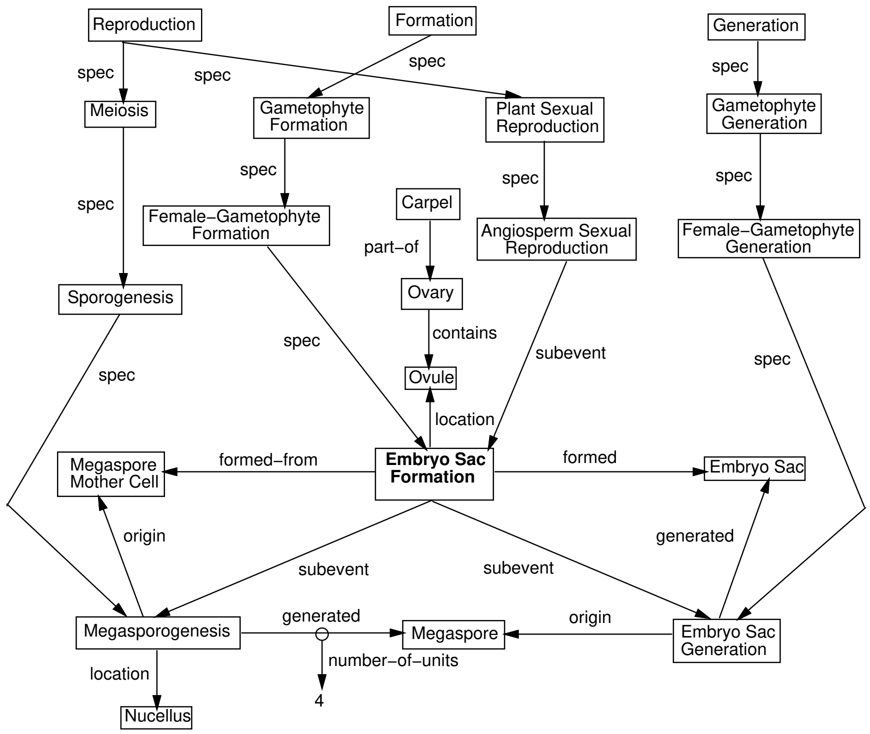

Figure 13: A representation of embryo sac formation

tent specifications, enabled Knight to construct two *Participants* views in the explanation plan: one for each of the subevents in *embryo sac formation* (Figure 13). When the views in this explanation plan are translated to text by the realization system, Knight produces the following text:

> Embryo sac formation is a kind of female gametophyte formation. During embryo sac formation, the embryo sac is formed from the megaspore mother cell. Embryo sac formation occurs in the ovule.
>
> Embryo sac formation is a step of angiosperm sexual reproduction. It consists of megasporogenesis and embryo sac generation. During megasporogenesis, the megaspore mother cell divides in the nucellus to form 4 megaspores. During embryo sac generation, the embryo sac is generated from the megaspore.

These algorithms have been used to generate explanations about hundreds of different concepts in the Biology Knowledge Base. For example, Section 2 shows other explanations generated by Knight. The explanation of *pollen tube growth* was produced by applying the *Explain-Process* EDP, and the explanations of *spore* and *root system* were produced by applying the *Explain-Object* EDP.

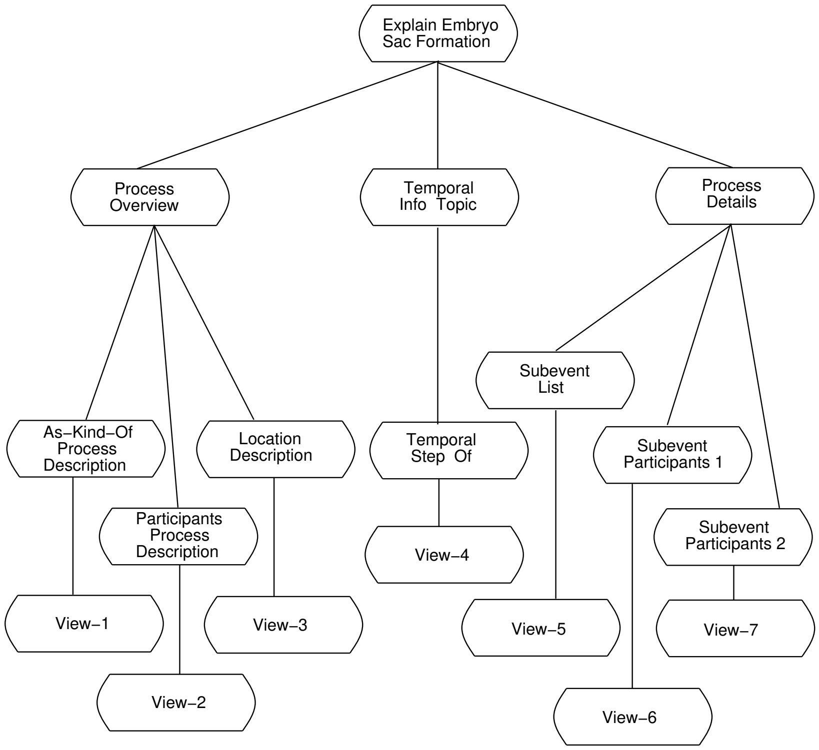

Figure 14: An explanation plan for embryo sac formation: High verbosity

8 Evaluation

Traditionally, research projects in explanation generation have not included empirical evaluations. Conducting a formal study with a generator has posed difficulties for at least three reasons: the absence of large-scale knowledge bases; the problem of robustness; and the subjective nature of the task. The field of explanation generation has experienced a dearth of "raw materials." The task of an explanation generator is three-fold: to extract information from a knowledge base, to organize this information, and to translate it to natural language. Unless an explanation generator has access to a sufficiently large knowledge base, the first step—and hence the second and third—cannot be carried out enough times to evaluate the system empirically. Unfortunately, because of the tremendous cost of construction, large-scale knowledge bases are scarce.

Even if large-scale knowledge bases were more plentiful, an explanation generator cannot be evaluated on them unless it was specifically designed to perform robustly. In very practical terms, a generator is likely to halt abruptly when it encounters unusual and unexpected knowledge structures; if this happens frequently, the system cannot be evaluated. We conjecture that most implemented explanation generators would meet with serious difficulties when applied to a large-scale knowledge base.

Furthermore, explanation generation is an ill-defined task. It stands in contrast to a machine learning task such as rule induction from examples. Although one can easily count the number of examples that an induction program classifies correctly, there is no corresponding objective metric for an explanation generator. Ideally, we would like to "measure" the coherence of explanations. Although it is clear that coherence is of paramount importance for explanation generation, there is no litmus test for it.

Given these difficulties, how can one evaluate the architectures, algorithms, and knowledge structures that form the basis for an explanation generator? The traditional approach has been to

1. Conduct an analytical evaluation of a system's architecture and principal,
2. Demonstrate that it can produce well-formed explanations on a few examples.

While these evaluation techniques are important, they are not sufficient. Three steps can be taken to enable better evaluation. First, we can construct large-scale knowledge bases, such as the Biology Knowledge Base. Second, we can design and implement robust explanation systems that employ a representation of discourse knowledge that is easily manipulable by discourse-knowledge engineers. Third, to ensure that a knowledge base is not tailored for the purposes of explanation generation, we can enter into a "contractual agreement" with knowledge engineers; this eliminates all requests for representational modifications that would skew the representation to the task of explanation generation.

8.1 Experimental Design

The **Two-Panel Evaluation Methodology** can be used to empirically evaluate natural language generation work. We developed this methodology, which involves two panels of domain experts, to combat the inherent subjectivity of NLG: although multiple judges will rarely reach a consensus, their collective opinion provides persuasive evidence about the quality of explanations. To ensure the integrity of the evaluation results, a central stipulation of the methodology is that the following condition be maintained throughout the study:

> **Computer Blindness:** None of the participants can be aware that some texts are machine-generated or, for that matter, that a computer is in any way involved in the study.

The methodology involves four steps:

1. Generation of explanations by computer.
2. Formation of two panels of domain experts.
3. Generation of explanations by one panel of domain experts.
4. Evaluation of all explanations by second panel of domain experts.

Each of these is discussed in turn.

Explanation Generation: Knight. Because KNIGHT's operation is initiated when a user poses a question, the first task was to select the questions it would be asked. To this end, we combed the Biology Knowledge Base for concepts that could furnish topics for questions. Although the knowledge base focuses on botanical anatomy, physiology, and development, it also contains a substantial amount of information about biological taxons. Because this latter area is significantly less developed, we ruled out concepts about taxons. In addition, we ruled out concepts that were too abstract, e.g., *Object*. We then requested KNIGHT to generate explanations about the 388 concepts that passed through these filters.

To thoroughly exercise KNIGHT's organizational abilities, we were most interested in observing its performance on longer explanations. Hence, we eliminated explanations of concepts that were sparsely represented in the knowledge base. To this end, we passed the 388 explanations through a "length filter": explanations that consisted of at least 3 sentences were retained; shorter explanations were disposed of. This produced 87 explanations, of which 48 described objects and 39 described processes. Finally, to test an equal number of objects and processes, we randomly chose 30 objects and 30 process. Sample explanations produced by KNIGHT are shown in Appendix A.

Two Panels of Domain Experts. To address the difficult problem of subjectivity, we assembled 12 domain experts, all of whom were PhD students and post-doctoral scientists in biology. Because we wanted to gauge KNIGHT's performance relative to humans, we assigned each of the experts to one of two panels: the *Writing Panel* and the *Judging Panel*. By securing the services of such a large number of domain experts, we were able to form relatively large panels of four writers and eight judges (Figure 15). To ensure that the human-generated explanations would be of the highest possible quality, we assigned the four most experienced experts to the Writing Panel. The remaining eight experts were assigned to the Judging Panel to evaluate explanations.

To minimize the effect of factors that might make it difficult for judges to compare KNIGHT's explanations with those of domain experts, we took three precautions. First, we attempted to control for the length of explanations. Although we could not impose hard constraints, we made suggestions about how long a typical explanation might be. Second, to make the "level" of the explanations comparable, we asked writers to compose explanations for a particular audience, freshman biology students. Third, so that the general topics of discussion would be comparable, we asked writers to focus on anatomy, physiology, and development.

Explanation Generation: Humans. To ensure that the difficulty of the concepts assigned to the writers were the same as those assigned to KNIGHT, the writers were given the task of explaining *exactly* the same set of concepts that KNIGHT had explained. Because we wanted to give writers an opportunity to explain both objects and processes, each writer was given an approximately equal number of objects and processes. Each of the 4 writers was given 15 concepts to explain, and each concept was assigned to exactly one writer. We then transcribed their handwritten explanations and put them and KNIGHT's explanations into an identical format. Sample explanations produced

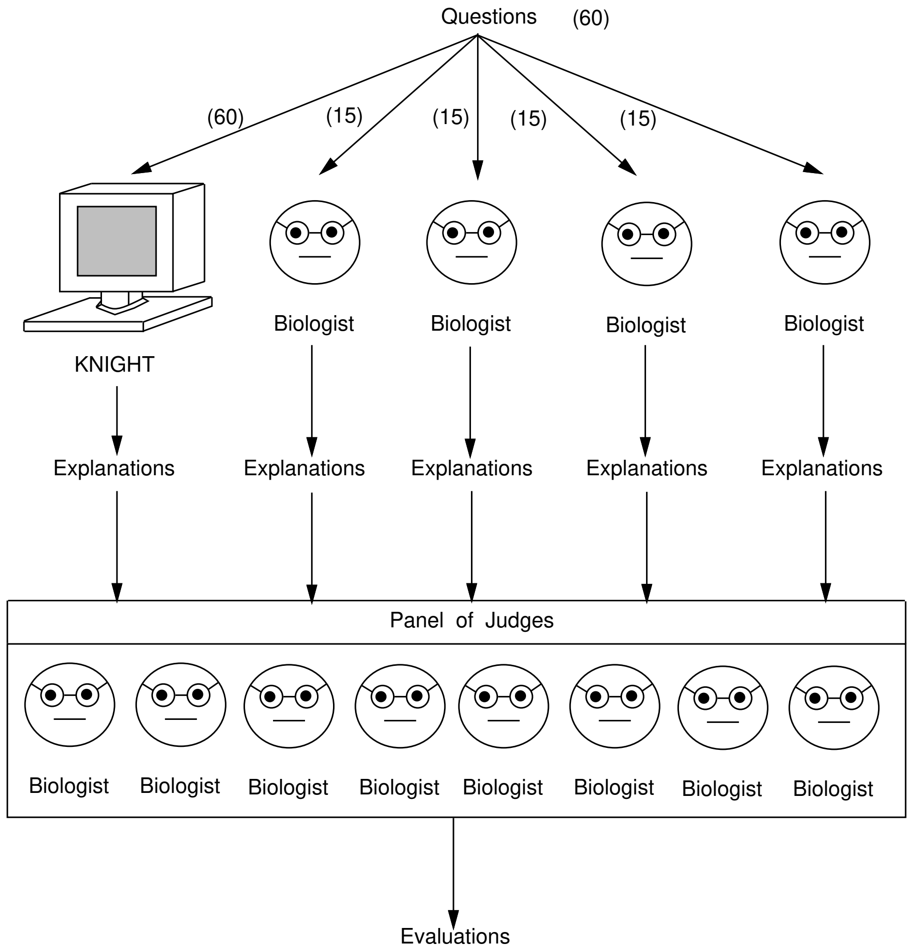

Figure 15: The Two-Panel Methodology in the KNIGHT Experiments

by the writers are shown in Appendix B. At this point, we had a pool of 120 explanations: sixty of these pertained to objects (30 written by biologists and 30 by KNIGHT), and the other sixty pertained to processes (also 30 written by biologists and 30 by KNIGHT).

Explanation Evaluation. We then submitted the explanations to the panel of eight judges. The judges were not informed of the source of the explanations, and all of the explanations appeared in the same format. Each judge was given fifteen explanations to evaluate. Judges were asked to rate the explanations on several dimensions: overall quality and coherence, content, organization, writing style, and correctness. To provide judges with a familiar rating scale, they were asked to assign letters grades (A, B, C, D, or F) to each explanation on each of the dimensions. Because carefully evaluating multiple dimensions of explanations is a labor-intensive task, time considerations required us to limit the number of explanations submitted to each judge. Hence, we assigned each judge 15 explanations, which on average required an hour to evaluate. We assigned explanations to judges using an allocation policy that obeyed the following four constraints:

- *System-Human Division:* Each judge received explanations that were approximately evenly divided between those that were produced by KNIGHT and those that were produced by biologists.
- *Object-Process Division:* Each judge received explanations that were approximately evenly divided between objects and processes.
- *Single-Explanation Restriction:* No judge received two explanations of the same concept.
- *Multi-Judge Stipulation:* The explanations written by each writer were not evaluated by only one judge; rather, they were distributed to at least two judges.

It is important to emphasize again that the judges were not made aware of the purpose of the experiment, nor were told that any of the explanations were computer-generated.

8.2 Results

By the end of the study, we had amassed a large volume of data. To analyze it, we converted each of the "grades" to their traditional numerical counterparts, i.e., A=4, B=3, C=2, D=1, and F=0. Next, we computed means and standard errors for both KNIGHT's and the biologists' grades. We calculated these values for the overall quality and coherence rating, as well as for each of the dimensions of content, organization, writing style, and correctness. On the overall rating and on each of the dimensions, KNIGHT scored within approximately "half a grade" of the biologists (Table 3).[15]

Given these results, we decided to investigate the differences between KNIGHT's grades and the biologists' grades. When we normalized the grades by defining an "A" to be the mean of the biologists' grades, KNIGHT earned approximately 3.5 (a B^+). Comparing differences in dimensions, KNIGHT performed best on correctness and content, not quite as well on writing style, and least well on organization.

Because the differences between KNIGHT and the biologists were narrow in some cases, we measured the statistical significance of these differences by running standard t-tests.[16] KNIGHT's grades on the content, organization, and correctness dimensions did not differ significantly from the biologists' (Table 4). Of course, an insignificant difference does not indicate that KNIGHT's performance and the biologists' performance was equivalent—an even larger sample size might have shown a significant difference—however, it serves as an indicator that KNIGHT's performance approaches that of the biologists on these three dimensions.

To gauge how well KNIGHT generates explanations about objects—as opposed to processes—we computed means and standard errors for both KNIGHT's explanations of objects and the biologists'

[15]In the tables, $\pm$ denotes the standard error, i.e., the standard deviation of the mean.

[16]All t-tests were unpaired, two-tailed. The results are reported for a 0.05 level of confidence.

explanations of objects. We did the same for the explanations of processes. For both objects and processes, KNIGHT scored within "half a grade" of the biologists. Again, we measured the statistical significance of these differences. Although there was a significant difference between KNIGHT and biologists on explanations of processes, KNIGHT and the biologists did not differ significantly on explanations of objects (Tables 5 and 6).

As a final test, we compared KNIGHT to each of the individual writers. For a given writer, we assessed KNIGHT's performance relative to that writer in the following way: we compared the grades awarded to KNIGHT and the grades awarded to the writer on explanations generated in response to the same set of questions. This analysis produced some surprising results. Although there were substantial differences between KNIGHT and "Writer 1," KNIGHT was somewhat closer to "Writer 2," it was very close to "Writer 3," and its performance actually exceeded that of "Writer 4." KNIGHT and Writers 2, 3, and 4 did not differ significantly (Table 7).

9 Related Work

By synthesizing a broad range of research in natural language generation, KNIGHT provides a "start-to-finish" solution to the problem of automatically constructing expository explanations from semantically rich, large-scale knowledge bases. It introduces a new evaluation methodology and builds on the conceptual framework that has evolved in the NLG community over the past decade, particularly in techniques for knowledge-base access and discourse knowledge representation. We discuss each of these in turn.

Evaluation Methodologies. With regard to evaluation, KNIGHT is perhaps most closely related to five NLG projects that have been empirically evaluated: PAULINE [23], EDGE [8], the Example Generator[17] [46], ANA [27], and STREAK [54]. By varying pragmatic information such as tone, Hovy enabled PAULINE to generate many different paragraphs on the same topic. PAULINE's texts were not formally analyzed by a panel of judges, and it did not produce texts on a wide range of topics (it generated texts on only three different events.) However, this project is a significant achievement in terms of evaluation *scale* because of the sheer number of texts it produced: PAULINE generated more than 100 different paragraphs on the same subject. In a second landmark evaluation, Cawsey undertook a study in which subjects were allowed to interact with her explanation generation system, EDGE [8]. Subjects posed questions to EDGE about the operation of four circuits. Cawsey analyzed the system's behavior as the dialogs progressed, interviewed subjects, and used the results to revise the system. Although EDGE does not include a realization system (other than simple templates) and it was not subjected to a tightly controlled, formal evaluation, it was sufficiently robust to be used interactively by eight subjects.

The EXAMPLE GENERATOR [46], ANA [27], and STREAK [54] were each subjected to formal (quantitative) evaluations. Mittal and Paris developed and formally evaluated a generator that

[17] Mittal and Paris' system has no official name; we refer to it as "the Example Generator" for ease of reference.

Generator	*Overall*	*Content*	*Organization*	*Writing*	*Correctness*
KNIGHT	2.37±0.13	2.65±0.13	2.45±0.16	2.40±0.13	3.07±0.15
Human	2.85±0.15	2.95±0.16	3.07±0.16	2.93±0.16	3.16±0.15

Table 3: Comprehensive Analysis

	Overall	*Content*	*Organization*	*Writing*	*Correctness*
Difference	0.48	0.30	0.62	0.53	0.09
t statistic	-2.36	-1.47	-2.73	-2.54	-0.42
Significance	0.02	0.14	0.07	0.01	0.67
Significant?	Yes	No	No	Yes	No

Table 4: Differences and Significance

Generator	*Grade*
KNIGHT	2.65±0.19
Human	2.93±0.19
Difference	0.28
t statistic	-1.05
Significance	0.30
Significant?	No

Table 5: Explanation of Objects

Generator	*Grade*
KNIGHT	2.10±0.24
Human	2.77±0.17
Difference	0.67
t statistic	-2.23
Significance	0.03
Significant?	Yes

Table 6: Explanation of Processes

produced descriptions integrating text and examples. Rather than evaluating the explanations directly, subjects were given a quiz about the concept under consideration.[18] The degree to which the experiments controlled for specific factors, e.g., the effect of example positioning, example types, example complexity, and example order, is remarkable. ANA and STREAK were both subjected to quantitative, corpus-based evaluations. Kukich employed a corpus-based methodology to judge the coverage of ANA's knowledge structures. STREAK, which constructs summaries of basketball games, is part of of a larger effort by J. Robin, K. McKeown, and their colleagues at Columbia and Bellcore to develop robust document generation systems [43, 40]. It was evaluated with a corpus-based study that produced estimates of STREAK's sub-language coverage, extensibility, and the overall effectiveness of its revision-based generation techniques. Although neither of these studies employed human judges to critique text quality, the rigor with which they were conducted has significantly raised the standards for evaluating generation systems.

The relationship between the KNIGHT evaluation and those of its predecessors is summarized in Table 8. KNIGHT, STREAK, and ANA were all evaluated formally, i.e., quantitatively, while PAULINE and EDGE were evaluated informally. The KNIGHT, EDGE, and EXAMPLE GENERATOR evaluations employed humans as judges, while the ANA and STREAK evaluations had "artificial judges" in the form of corpora, and PAULINE was evaluated without judges. KNIGHT is the only system to have been evaluated in the context of a semantically rich, large-scale knowledge base. KNIGHT is also the only system to have been evaluated in a kind of restricted "Turing test" in which the quality of its text was evaluated by humans in a head-to-head comparison against the text produced by humans (domain experts) in response to the same set of questions.

Knowledge-Base Access. KNIGHT builds on a rich body of research on knowledge base access by view selection. Many different approaches have been taken to selecting views. In the knowledge representation community, views of a concept have been extracted by changing its superclass, as in

[18] In a second analysis *without human judges*, the system developers compared selected features of the EXAMPLE GENERATOR's output with text from textbook and obtained encouraging results.

Knight	*vs. Writer 1*	*vs. Writer 2*	*vs. Writer 3*	*vs. Writer 4*
Knight	1.93±0.29	2.73±0.23	2.73±0.27	2.07±0.23
Human	3.60±0.16	3.40±0.23	2.80±0.28	1.60±0.23
Difference	1.67	0.67	0.07	0.47
t statistic	-5.16	-2.03	-0.17	1.42
Significance	0.00	0.05	0.86	0.16
Significant?	Yes	No	No	No

Table 7: Knight vs. Individual Writers

	Pauline	Edge	Example Generator	Ana	Streak	Knight
Formality	Informal	Informal	Formal	Formal	Formal	Formal
"Judges"	None	Humans	Humans	Corpus	Corpus	Humans
Large-Scale KB?	No	No	No	No	No	Yes
System *vs.* Human	No	No	No*	No	No	Yes

Table 8: Evaluation Methodologies

the KRL [5] and KODIAK [66] systems. More recent work in this area has demonstrated how views can be "reified," i.e., encoded as first-class objects in their representations [57, 11, 28]. Related research in intelligent tutoring systems explored how different views could be used to model different aspects of the domain and, hence, to answer different kinds of questions [58]. Work in automated knowledge acquisition systems has shown how to use views to constrain search [50].

Several projects in explanation generation have exploited views to improve the quality of the explanations they provide. The ADVISOR system [44] represents views with a multiple-hierarchy knowledge base. ADVISOR infers a user's current goal from his most recent utterances and uses this goal to select a hierarchy from the multiple-hierarchy knowledge base. The selected view controls the content of the explanation and the reasoning that produced that content. In a similar vein, viewpoints in Swartout's XPLAIN [64] are annotations that indicate when to include a piece of knowledge in an explanation.

It is preferable to construct (i.e., extract) views at runtime rather than encoding them in a knowledge base. If a KB accessing system could dynamically construct views, the discourse-knowledge engineer would be freed from the task of having to anticipate all queries and rhetorical situations and precompiling semantic units for each situation. KNIGHT, ROMPER [39], and Suthers work [59, 60, 61, 63] use these types of views to determine the content of their explanations. Once a perspective is selected, ROMPER includes in its explanations only those attributes whose salience values are the highest. In contrast to ROMPER's views which are domain-specific, i.e., its views are confined to the domain of financial securities, Suthers' and KNIGHT's views are domain independent. Suthers set forth a set of views which can be used to select coherent subsets of domain knowledge: structural, functional, causal, constraint, and process. He also developed a view retriever and a highly refined theory of explanation generation in which views play a significant role. KNIGHT's views are very similar to McCoy's and Suthers in that they define the relations and properties of a concept that are relevant when considering a concept from a viewpoint belonging to that view type [3, 1]. They also provide four types of knowledge-base access robustness, as discussed in Section 3.

Discourse Generation. Two principle mechanisms have been developed for generating discourse: schemata and top-down planners.[19] McKeown's pioneering work on *schemata* marks the beginning of the "modern era" of discourse generation [41]. Schemata are ATN-like structures that represent naturally occurring patterns of discourse. For example, a schema for defining a concept includes instructions to identify its superclass, to name its parts, and to list its attributes. To construct an explanation plan, McKeown's TEXT system traverses the schemata and sequentially instantiates rhetorical predicates with propositions from a knowledge base. Paris extended schemata to generate descriptions of complex objects in a manner that is appropriate for the user's level of expertise [52], and ROMPER's schemata include information about the content of proposi-

[19]A third alternative proposed by Sibun are *short-range* strategies that exploit relations such as spatial proximity to guide the generator through the domain knowledge [55]. These strategies are designed for flexibility, but they do not account for extended explanations, which require a more global rhetorical structure.

tions to be selected, as well as their communicative role.[20] Although schemata have been criticized because they lack flexibility, they successfully capture many aspects of discourse structure.

An alternative to schemata is the top-down planning approach [49, 62, 8, 37, 24, 48, 47].[21] The operators of two of these planning systems are based on Rhetorical Structure Theory (RST) [35]. Hovy's Structurer [22, 21, 24] is a hierarchical planner whose operators instantiate relations from RST, and Hovy and McCoy have developed methods for increasing the organizational flexibility of the Structurer [25]. The Reactive Planner also uses RST-like operators. However, unlike all of the preceding research—and unlike KNIGHT as well—it offers sophisticated mechanisms for generating explanations in interactive contexts [49, 48, 47]. Because the operators explicitly record the rhetorical effects achieved, and because the system records alternative operators it could have chosen, as well as assumptions it made about the user, the Reactive Planner can respond to follow-up questions—even if they are ambiguous—in a principled manner. A related approach has been taken by Cawsey in the EDGE system [8]. Because EDGE has facilities for managing conversations, users may interrupt the system to ask questions, and EDGE can either answer the question immediately or postpone its response. Suthers [62] has developed a sophisticated hybrid approach that includes planning techniques as well as plan critics, simulation models, reorganization methods, and graph traversal. By assembling these diverse mechanisms into a single architecture, he demonstrates how the complexities of explanation planning can be dealt with in a coherent framework. The principal advantage of top-down planners over schema-based generators is their ability to reason about the structure, content, and goals of explanations—as opposed to merely instantiating pre-existing plans embodied by schemata.

KNIGHT's EDPs are much more schema-like than plan-like. Although EDPs have inclusion conditions, which are similar to the *constraint* attribute of RST-based operators, and they provide a *centrality* attribute, which enables KNIGHT to reason about the inclusion of a topic if "space" is limited, EDPs do not in general permit KNIGHT to reason about the goals fulfilled by particular text segments as do plan-based systems. However, EDPs have proven to be very useful as a discourse-knowledge engineering tool, a result that can be attributed in large part to their combining a frame-based representation with procedural constructs. In a sense, EDPs are schemata whose representation has been fine-tuned to maximize ease of use on a large scale.

10 Future Directions

Research on KNIGHT suggests several directions for future work. First, the results of the evaluation call for further analysis and experimentation. For example, an in-depth analysis at the discourse, sentential and lexical levels of all of the texts produced by both the humans and the system may

[20] Unlike McKeown's and Paris' schemata—as well as KNIGHT's EDPs—ROMPER's schemata do not impose an order on the propositions.

[21] The planning approach, which has dominated the field for the past few years can be traced to Appelt's work on planning referring expressions [4], which itself builds on earlier work on reasoning about speech acts in a planning paradigm [9].

reveal which characteristics of the highly rated texts are desirable. These in turn can be used to improve the EDPs. On the empirical side, a particularly intriguing kind of experiment is an *ablation* study. In an ablation study, different aspects of the system are ablated (removed or degraded), and the effects of the ablation on the explanations are noted. By performing a series of these experiments, one can determine which aspects of the system and its representational structures contribute most significantly to its success.

Second, although EDPs were employed successfully in the generation of hundreds of explanations, the fact that they have more in common with schemata than with the operators of top-down planners is indicative of a fundamental limitation: the intentional structure of the discourse is unavailable for inference. As a result, it is considerably more difficult for the system to respond to follow-up questions, reason about paragraph structure, perform goal-based content determination, and produce discourse cues. This calls for the incorporation of an intentional structure into EDPs, but modifying EDPs to represent intention must be accomplished in such a way that the "discourse-knowledge engineering" properties are preserved and there is no sacrifice of text quality. Moreover, KNIGHT currently employs very rudimentary pronominalization techniques. Including more sophisticated methods [12] should result in a significant increase in text quality.

Finally, one of the most fruitful area for future work is research on multimedia explanation generation [17, 36, 65, 51, 10]. A multimedia explanation generator can combine the explanatory capabilities of KNIGHT with the powerful display technologies of multimedia devices. By tightly coupling multimedia object indices to a formal representation of domain knowledge, a multimedia explanation generator should be able to dynamically generate responses involving text, graphics, animation, audio, and video. To this end, we have begun work on an architecture for multimedia explanation generation [34]. This kind of system must be able to address issues of multimedia content determination and organization, as well as determining at runtime which media should be used to realize the concepts to be communicated.

11 Conclusion

Explanation generation is an exceedingly complex task that involves a diversity of interacting computational mechanisms. An explanation system must be able to select from a knowledge base precisely those facts that enable it to construct a response to a user's question, organize this information and translate the formal representational structures found in knowledge bases to natural language. While the traditional approach to work on explanation has been to develop a proof-of-concept system and to demonstrate that it can produce well-formed explanations on a few examples, developing robust explanation generation techniques and scalable discourse knowledge representations facilitates more extensive, empirical studies.

To investigate the issues and problems of generating natural language explanations from semantically rich, large-scale knowledge bases, we have designed and implemented KNIGHT, a fully functioning explanation system that automatically constructs multi-sentential and multi-paragraph natural language explanations. KNIGHT has generated hundreds of explanations from the Biology Knowledge Base. It addresses a multiplicity of issues in explanation generation, ranging from knowledge base access and discourse planning to a new methodology for empirical evaluation. This work has demonstrated that (1) separating out knowledge-base access from explanation planning can enable the construction of a robust system that extracts coherent views from a a semantically rich, large-scale knowledge base; and (2) Explanation Design Packages, a hybrid representation of discourse knowledge that combines a frame-based representation with procedural constructs, facilitate the iterative refinement of discourse knowledge. Combining hierarchically structured discourse objects with embedded procedural constructs, EDPs have been used to represent discourse knowledge about explaining physical objects and processes, and they have been tested in the generation of hundreds of explanations of biological concepts.

To gauge the effectiveness of these techniques, we developed the Two-Panel Evaluation Methodology and employed it in the evaluation of KNIGHT. KNIGHT scored within "half a grade" of the biologists. There was no significant difference between KNIGHT's explanations and the biologists' explanations on measures of content, organization, and correctness, nor was there a statistically significant difference in overall quality between KNIGHT's explanations those composed by three of the biologists. KNIGHT's performance exceeded that of one of the biologists.

In summary, it is encouraging that an explanation system could begin to approach the performance of multiple domain experts and surpass that of one expert. These findings demonstrate that an explanation system that has been given a well represented knowledge base can construct natural language responses whose quality approximates that of humans. More generally, they suggest that we are beginning to witness the appearance of computational machinery that will significantly broaden the bandwidth of human-computer communication.

Acknowledgements

We would like to thank our principle domain expert, Art Souther, for leading the knowledge base construction effort; Charles Callaway and the NLG students for their work on the realization system; Erik Eilerts, for building the knowledge base editing tools; Michael Elhadad, for generously assisting us with FUF; Peter Clark and Charles Callaway for helpful comments on previous drafts of this paper; Dan Suthers for insights on the problems of evaluating explanation systems; and the members of the Biology Knowledge Base Project: Liane Acker, Brad Blumenthal, Rich Mallory, Ken Murray, and Jeff Rickel.

References

[1] L. E. H. Acker. *Access Methods for Large, Multifunctional Knowledge Bases.* PhD thesis, The University of Texas at Austin, Austin, Texas, 1992.

[2] L. E. H. Acker and B. W. Porter. Extracting viewpoints from knowledge bases. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*, pages 547–552, 1994.

[3] L. H. Acker, J. C. Lester, A. F. Souther, and B. W. Porter. Generating coherent explanations to answer students' questions. In H. Burns, J. Parlett, and C. Redfield, editors, *Intelligent Tutoring Systems: Evolutions in Design*, pages 151–176. Lawrence Earlbaum, Hillsdale, New Jersey, 1991.

[4] D. E. Appelt. Planning english referring expressions. *Artificial Intelligence*, 26:1–33, 1985.

[5] D. G. Bobrow and T. Winograd. An overview of KRL, a knowledge representation language. *Cognitive Science*, 1:3–46, 1977.

[6] B. G. Buchanan and E. H. Shortliffe, editors. *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project.* Addison-Wesley, Reading, Massachusetts, 1984.

[7] C. B. Callaway and J. C. Lester. Robust natural language generation from large-scale knowledge bases. In *Proceedings of the Fourth Bar-Ilan Symposium on Foundations of Artificial Intelligence*, pages 96–105, 1995.

[8] A. Cawsey. *Explanation and Interaction: The Computer Generation of Explanatory Dialogues.* MIT Press, 1992.

[9] P. R. Cohen and C. R. Perrault. Elements of a plan-based theory of speech acts. *Cognitive Science*, 3:177–212, 1979.

[10] M. Cornell, D. Suthers, and B. Woolf. Using "live information" in a multimedia framework. In M. T. Maybury, editor, *Intelligent Multimedia Interfaces*, chapter 13. AAAI Press, 1993.

[11] J. Crawford. Access-limited logic—a language for knowledge representation. Technical Report AI Laboratory AI90-141, University of Texas at Austin, Austin, Texas, 1990.

[12] R. Dale. *Generating Referring Expressions.* MIT Press, 1992.

[13] R. de Beaugrande. *Text Discourse and Process.* Ablex, Norwood, NJ, 1980.

[14] E. Eilerts. KnEd: An interface for a frame-based knowledge representation system. Master's thesis, The University of Texas at Austin, Austin, Texas, May 1994.

[15] M. Elhadad. FUF: The universal unifier user manual version 5.0. Technical Report CUCS-038-91, Department of Computer Science, Columbia University, 1991.

[16] M. Elhadad. *Using Argumentation to Control Lexical Choice: A Functional Unification Implementation.* PhD thesis, Columbia University, 1992.

[17] S. K. Feiner and K. R. McKeown. Coordinating text and graphics in explanation generation. In *Proceedings of the Eighth National Conference on Artificial Intelligence*, pages 442–449, Boston, Massachusetts, 1990.

[18] J. E. Grimes. *The Thread of Discourse.* Mouton, The Hague, 1975.

[19] M. A. K. Halliday and R. Hassan. *Cohesion in English.* Longman, London, 1976.

[20] J. R. Hobbs. On the coherence and structure of discourse. Technical Report CSLI-85-37, Center for the Study of Language and Information, Stanford University, Stanford, California, October 1985.

[21] E. H. Hovy. Approaches to the planning of coherent text. In *Fourth International Workshop on Text Generation*, Catalina Island, California, July 1988.

[22] E. H. Hovy. Planning coherent multisentential text. In *Proceedings of the 26th Annual Meeting*, pages 163–169, Buffalo, 1988. Association for Computational Linguistics.

[23] E. H. Hovy. Pragmatics and natural language generation. *Artificial Intelligence*, 43:153–197, 1990.

[24] E. H. Hovy. Automated discourse generation using discourse structure relations. *Artificial Intelligence*, 63:341–385, 1993.

[25] E. H. Hovy and K. F. McCoy. Focusing your RST: A step toward generating coherent multisentential text. In *Proceedings of the Eleventh Cognitive Science Society Conference*, pages 667–674, Ann Arbor, Michigan, 1989.

[26] A. K. Joshi and S. Weinstein. Control of inference: The role of some aspects of discourse structure—centering. In *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, pages 385–387, Cambridge, England, 1981.

[27] K. Kukich. *Knowledge-Based Report Generation: A Knowledge Engineering Approach to Natural Language Report Generation.* PhD thesis, University of Pittsburgh, 1983.

[28] D. B. Lenat and R. V. Guha. *Building Large Knowledge Based Systems.* Addison-Wesley, Reading, Massachusetts, 1990.

[29] J. Lester. *Generating Natural Language Explanations from Large-Scale Knowledge Bases.* PhD thesis, The University of Texas at Austin, Austin, Texas, 1994.

[30] J. C. Lester and B. W. Porter. Generating integrative explanations: A delayed-commitment approach. In *Proceedings of the AAAI Workshop on Explanation*, pages 80–89, Boston, Massachusetts, July 1990.

[31] J. C. Lester and B. W. Porter. An architecture for planning multi-paragraph pedagogical explanations. In *AAAI Workshop on Comparative Analysis of Explanation Planning Architectures*, pages 27–41, 1991.

[32] J. C. Lester and B. W. Porter. A revision-based model of instructional multi-paragraph discourse production. In *Proceedings of the Thirteenth Cognitive Science Society Conference*, pages 796–800, 1991.

[33] J. C. Lester and B. W. Porter. A student-sensitive discourse generator for intelligent tutoring systems. In *Proceedings of the International Conference on the Learning Sciences*, pages 298–304, August 1991.

[34] J. C. Lester and B. W. Porter. Designing multi-media knowledge delivery systems: The strong representation paradigm. In *Working Notes of the AAAI Spring Symposium on Intelligent Multi-Media Multi-Modal Systems*, 1994.

[35] W. C. Mann and S. A. Thompson. Rhetorical structure theory: A theory of text organization. Technical Report ISI/RS-87-190, USC/Information Sciences Institute, Marina del Rey, California, June 1987.

[36] M. T. Maybury. Planning multimedia explanations using communicative acts. In *Proceedings of the Ninth National Conference on Artificial Intelligence*, pages 61–66, Anaheim, California, 1991.

[37] M. T. Maybury. Communicative acts for explanation generation. *International Journal of Man-Machine Studies*, 37(2):135–172, 1992.

[38] M. T. Maybury. Communicative acts for generating natural language arguments. In *Proceedings of the Eleventh National Conference on Artificial Intelligence*, pages 357–364, 1993.

[39] K. F. McCoy. Generating context-sensitive responses to object-related misconceptions. *Artificial Intelligence*, 41:157–195, 1989 1990.

[40] K. McKeown, J. Robin, and K. Kukick. Generating concise natural language summaries. *Information Processing and Management*, 1995. Special Issue on Summarization.

[41] K. R. McKeown. *Generating Natural Language Text in Response to Questions about Database Structure.* PhD thesis, University of Pennsylvania, 1982.

[42] K. R. McKeown. *Text Generation: Using Discourse Strategies and Focus Constraints to Generate Natural Language Text.* Cambridge University Press, 1985.

[43] K. R. McKeown, K. Kukich, and J. Shaw. Practical issues in automatic documentation generation. In *Proceedings of the Conference on Applied Natural Language Processing*, 1994.

[44] K. R. McKeown, M. Wish, and K. Matthews. Tailoring explanations for the user. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 794–798, 1985.

[45] K. P. Mitchell. A system to generate natural language sentences for explanations from a large knowledge base. Master's thesis, The University of Texas at Austin, Austin, Texas, May 1992.

[46] V. O. Mittal. *Generating Natural Language Descriptions with Integrated Text and Examples.* PhD thesis, University of Southern California, September 1993.

[47] J. D. Moore. *Participating in Explanatory Dialogues.* MIT Press, 1995.

[48] J. D. Moore and C. L. Paris. Planning text for advisory dialogues: Capturing intentional and rhetorical information. *Computational Linguistics*, 19(4):651–694, 1993.

[49] J. D. Moore and W. R. Swartout. A reactive approach to explanation: Taking the user's feedback into account. In C. L. Paris, W. R. Swartout, and W. C. Mann, editors, *Natural Language Generation in Artificial Intelligence and Computational Linguistics*, pages 3–48. Kluwer Academic Publishers, Boston, 1991.

[50] K. S. Murray and B. W. Porter. Controlling search for the consequences of new information during knowledge integration. In *Proceedings of the Sixth International Workshop on Machine Learning*, pages 290–295, 1989.

[51] J. G. Neal and S. C. Shapiro. Intelligent multi-media interface technology. In J. W. Sullivan and S. W. Tyler, editors, *Intelligent User Interfaces*, pages 11–43. Addison-Wesley, New York, 1991.

[52] C. L. Paris. Tailoring object descriptions to a user's level of expertise. *Computational Linguistics*, 14(3):64–78, September 1988.

[53] B. Porter, J. Lester, K. Murray, K. Pittman, A. Souther, L. Acker, and T. Jones. AI research in the context of a multifunctional knowledge base: The botany knowledge base project. Technical Report AI Laboratory AI88-88, University of Texas at Austin, Austin, Texas, 1988.

[54] J. Robin. *Revision-Based Generation of Natural Language Summaries Providing Historical Background.* PhD thesis, Columbia University, December 1994.

[55] P. Sibun. Generating text without trees. *Computational Intelligence*, 8(1):102–122, March 1992.

[56] A. Souther, L. Acker, J. Lester, and B. Porter. Using view types to generate explanations in intelligent tutoring systems. In *Proceedings of the Eleventh Cognitive Science Society Conference*, pages 123–130, 1989.

[57] J. F. Sowa. *Conceptual Structures.* Addison-Wesley, Reading, Massachusetts, 1984.

[58] A. L. Stevens and A. Collins. Multiple conceptual models of a complex system. In R. E. Snow, P. Federico, and W. E. Montague, editors, *Aptitutde, Learning, and Instruction*, volume 2, pages 177–197. Lawrence Earlbaum, Hillsdale, New Jersey, 1980.

[59] D. D. Suthers. Providing multiple views for explanation. In *Proceedings of the AAAI-88 Workshop on Explanation*, pages 12–15, 1988.

[60] D. D. Suthers. Perspectives in explanation. Technical Report COINS 89-24, University of Massachusetts, Amherst, Massachusetts, 1989.

[61] D. D. Suthers. The epistemological structure of explanation. In *Proceedings of the AAAI-90 Workshop on Explanation*, pages 178–187, 1990.

[62] D. D. Suthers. A task-appropriate hybrid architecture for explanation. *Computational Intelligence*, 7(4):315–333, November 1991.

[63] D. D. Suthers. *An Analysis of Explanation and Its Implications for the Design of Explanation Planners.* PhD thesis, University of Massachusetts, February 1993.

[64] W. R. Swartout. XPLAIN: A system for creating and explaining expert consulting programs. *Artificial Intelligence*, 21:285–325, 1983.

[65] W. Wahlster, E. André, W. Finkler, H.-J. Profitlich, and T. Rist. Plan-based integration of natural language and graphics generation. *Artificial Intelligence*, 63:387–427, 1993.

[66] R. Wilensky. KODIAK—a knowledge representation language. In *Proceedings of the Sixth AnnualConference of the Cognitive Science Society*, pages 344–352, 1984.

Appendix A: Explanations Generated by Knight

This appendix presents fifteen sample explanations generated by KNIGHT from the Biology Knowledge Base. These explanations were among the sixty explanations generated by KNIGHT that were judged by domain experts in the empirical evaluation.

```
#1    Fascicular Cambium

      Fascicular cambium is a kind of vascular cambium.  It is a
      component of the vascular bundle and is lateral to the xylem.
      It is medial to the phloem and is surrounded by the vascular
      bundle sheath.

#2    Marginal Meristem

      Marginal meristem is a subregion of the leaf.  It gives rise to
      plate meristem.  Plate meristem generation is a step of plate
      meristem development.  Marginal meristem gives rise to leaf
      protoderm.  Leaf protoderm generation is a step of plate
      meristem generation.  Marginal meristem gives rise to leaf
      procambium.  Leaf procambium generation is a step of plate
      meristem generation.  Marginal meristem gives rise to leaf
      promesophyll.  Leaf promesophyll generation is a step of plate
      meristem generation.

      Marginal meristem is composed of marginal initials and
      submarginal initials.

#3    Leaf Promesophyll

      Leaf promesophyll is a kind of immature tissue.  It is a
      subregion of plate meristem.  It is generated from marginal
      meristem.

#4    Endosperm Mother Cell

      The endosperm mother cell is a kind of female gamete of a plant.
      It is contained in the embryo sac.  It and the egg cell are
      produced from the megaspore during female gametogenesis.  It is
      fertilized by an angiosperm sperm cell to form the primary
      endosperm cell during endosperm fertilization.  Endosperm
      fertilization is a step of double fertilization.

      The endosperm mother cell contains 2 nuclei.
```

#5 Phloem

Phloem is a kind of vascular tissue. It is a component of the vascular bundle and is lateral to the fascicular cambium. It is surrounded by the vascular bundle sheath. It is composed of parenchyma, sclerenchyma and sieve tubes.

#6 Leaf Primordium

Leaf primordium is a kind of meristem. It is a subregion of the stem meristem and is below the shoot apical meristem. It is lateral to the sub-apical meristem. It is generated from shoot apical meristem during leaf initiation.

#7 Microspore

2 microspores are produced from the microspore mother cell during microsporogenesis. The microspore divides to form the pollen tube cell and the pollen generative cell during male gametogenesis. Male gametogenesis is a step of pollen grain development.

#8 Spore Mother Cell

The spore mother cell is a kind of diploid cell. It is part of the reproductive stage of a sporophyte. It divides to form 4 spores during sporogenesis. Sporogenesis is a step of gametophyte formation.

#9 Protoderm

Protoderm is a kind of primary meristem. It is a subregion of sub-apical meristem and surrounds ground meristem. It is generated from apical meristem.

#10 Plant Female Gamete

The female gamete of a plant is a kind of plant gamete. It is fertilized by the plant sperm cell to form the zygote during plant reproductive fertilization. Plant reproductive fertilization is a step of plant sexual reproduction.

#11 Corolla

The corolla is part of the flower and surrounds the androecium. It is surrounded by the calyx. It is made up of several petals.

#12 Embryo Sac

The embryo sac is a kind of female gametophyte. It is contained in the nucellus. It is generated from the megaspore. It contains the endosperm mother cell, 3 antipodal cells and the egg apparatus.

#13 Central Mother Cells

Central mother cells are a kind of immature tissue. They are a subregion of shoot apical meristem and are medial to the tunica. They are below the tunica and the apical initials.

#14 Shoot Apical Meristem

Shoot apical meristem is a subregion of the stem meristem and is above the leaf primordium and the sub-apical meristem. It gives rise to the leaf buttress during leaf buttress initiation. Leaf buttress initiation is a step of leaf initiation. Leaf primordium is generated from shoot apical meristem during leaf initiation. Leaf initiation is a step of leaf formation.

The subregions of the shoot apical meristem include the tunica, the central mother cells and the apical initials.

#15 Endodermis

The endodermis is a kind of ground tissue. It surrounds the pericycle. It is made up of endodermal cells. The parts of the endodermis include the cell walls, whose parts include the radial walls and the lateral walls, and the plasma membrane.

Appendix B: Explanations Generated by Domain Experts

This appendix presents fifteen sample explanations that were generated by domain experts. These explanations were among the sixty explanations generated by domain experts that were judged by a second panel of domain experts in the empirical evaluation. They pertain to the same topics as those in the preceding appendix.

```
#1    Fascicular Cambium

      the meristematic layer of cells within the vascular bundles that
      divides to produce secondary xylem and phloem.

#2    Marginal Meristem

      tissue with the capacity to divide mitotically, found in areas
      other than the apex of the plant, such as in the axials of
      leaves.  This tissue can produce new leaves, branches, or
      flowers depending on the hormonal signals received.

#3    Leaf Promesophyll

      in the developing leaf, the tissue that will differentiate into
      the two layers of the mesophyll, the palisade and spongy
      mesophyll.

#4    Endosperm Mother Cell

      the binucellate cell of the megagametophyte (female) that will
      be fertilized with the other sperm cell.  This triploid cell
      will then divide mitotically numerous times to produce endosperm
      tissue, which will be used by the developing embryo as a food
      source.

#5    Phloem

      the tissue within the vascular system that carries food (sugars)
      produced in the leaves by photosynthesis to the rest of the plant.
      Phloem cells are living when mature and functioning.

#6    Leaf Primordium

      the area of tissue within the meristem with the capacity to
      develop into a leaf.
```

#7 Microspore

the immature haploid gametophyte that will produce sperm. Produced by meiosis from a microspore mother cell within the anthers of the flower, the microspore will be protected by a hard resistant wall of sporopollenin to produce a pollen grain. The pollen grain germinates on the stigma and a microgametophyte develops.

#8 Spore Mother Cell

in algae, mosses, ferns, and bryophytes that produce only one type of gametophyte, the diploid cell that divides by meiosis to produce (usually) four haploid spores.

#9 Protoderm

The innermost cell layer of the epidermis of woody vascular plants. The protoderm is a site of epidermal cell division and differentiation.

#10 Plant Female Gamete

Haploid reproductive cell (often an egg) produced by the female gametophyte by mitosis. The female gamete is most often immotile and may be surrounded by other specialized haploid cells. The female gamete will fuse with the sperm (male gamete) to produce a diploid embryo.

#11 Corolla

The second outermost structure found in flowers. The corolla is surrounded to the outside by the calyx and to the inside by the androecium. The corolla may be brightly colored and showy or it may be extremely reduced or it may be missing entirely.

#12 Embryo Sac

The seven celled (eight nuclei) female gametophyte of angiosperms (flowering plants). The embryo sac consists of 1 egg cell surrounded by two companion cells, 3 antipodal cells, and a large central cell which contains 2 polar nuclei. The embryo sac is found inside the ovule of the sporophyte (diploid plant).

#13 Central Mother Cells

Diploid cells found within reproductive structures of plants that will undergo meiosis to produce four haploid cells known as spores. The spores may develop into a haploid plant known as the gametophyte.

#14 Shoot Apical Meristem

The growing tip of a stem or shoot. The apical meristem is the site of rapid cell division. Cell elongation and differentiation occur directly below the apical meristem.

#15 Endodermis

A cell layer within the root of vascular plants. The functions of the endodermis are to give rise to lateral roots and to assist in water movement within the root.