

Multimodal Reasoning in Molecular Imagery

Janice Glasgow¹ and Suzanne Fortier^{1,2}

Departments of ¹Computing and Information Science and ²Chemistry

Queen's University, Kingston, Canada

janice@qucis.queensu.ca fortiers@qucdn.queensu.ca

Abstract

Pattern recognition, model building, and image segmentation and interpretation are integral processes in the recovery of the three-dimensional structure of a protein model from experimentally derived data. The tools and techniques developed for computational imagery provide a multimodal framework for representing and reasoning with visual/spatial characteristics in this domain. This paper overviews the three distinct representations for computational imagery and describe their application to the problem of molecular scene analysis.

Introduction

How we represent and reason with molecular information depends on the type of questions we seek to answer. By making particular features of a molecule explicit, a representation can contain implicitly the protocols needed to answer specific queries. These representations may be very simple, such as the lists used to describe the primary structure of proteins, or they may take more complex forms, such as the projections used to depict configurational arrangements. They may be numerical, symbolic or visual, and they can encode information through spatial relationships (see Figure 1).

This paper describes a multimodal knowledge representation scheme for computational imagery and its application to *molecular scene analysis*. The concept of *scene analysis* has previously been used in the context of machine vision to refer to the set of processes associated with the reconstruction and classification of complex images. Such analyses rely on the availability of *a priori* domain information, in the form of templates or rules, to locate and identify features in the scene. By analogy, we use the phrase molecular scene analysis to refer to the processes associated with the interpretation and understanding of protein structures and molecular interactions.

An effective knowledge representation scheme must not only embody the salient features of relevant do-

main information as clearly and simply as possible, but it must also allow for efficient operations on this information. In the case of information pertaining to molecular structures, these can be divided into three broad classes: numerical, logical and imagery operations. Numerical operations are routinely implemented as computer algorithms. Logical and imagery operations, by contrast, are still mainly performed by experts. While most programs use simple logical functions, there are currently few implementations in the molecular domain that incorporate deductive reasoning formalisms. Similarly, graphics systems can transform and depict molecular images efficiently, but it is still the expert who performs the essential tasks of imagery, such as pattern recognition and spatial reasoning, that are associated with image understanding.

Research in computational imagery has focussed on automated tools for emulating the high-level processes associated with mental imagery, that is, the processes of constructing, transforming and inspecting image representations. The proposed knowledge representation framework for molecular imagery consists of three interrelated representations: a *descriptive* representation, which stores all the relevant information about a molecular domain and can be used for logical and numerical operations, and two *depictive* representations, which are used to carry out the visual and spatial analysis of an image. Our goal is to provide a framework within which all available knowledge, whether numerical, symbolic or depictive in nature, may be used to reason with and learn about molecular structures at varying levels of complexity. From this we can attempt to mimic the expert's deductive, experiential and imagery capabilities in solving the complex problems associated with a molecular scene analysis.

Section 2 of the paper presents a brief overview of molecular scene analysis. In Section 3 we demonstrate how the multimodal framework developed for computational imagery can be used to represent and reason with both descriptive and depictive knowledge in this

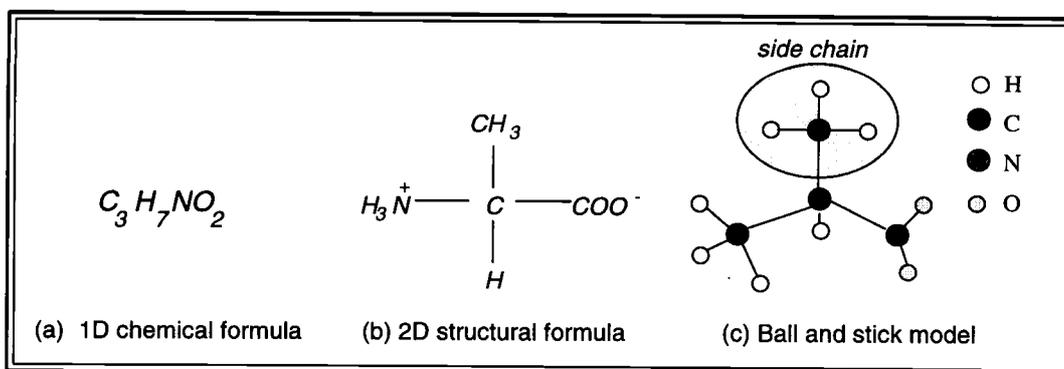


Figure 1: Alternative representations for the amino acid alanine

domain.

Molecular Scene Analysis

Research in machine vision has long been concerned with the problems involved in automatic image interpretation. Marr (Marr 1982) defines computational vision as the “*process of discovering what is present in the world, and where it is*”. Similar to the concept of scene analysis, which has emerged from research in machine vision, molecular scene analysis is concerned with the processes involved in the understanding of complex images of molecular structure.

We have previously presented a design for a computational approach to molecular scene analysis (Fortier *et al.* 1993; Glasgow, Fortier, & Allen 1993). The objective of this ongoing work is the implementation and application of scene analysis to the interpretation of protein structure. In the proposed approach, an electron density map (three-dimensional image) of a crystal is segmented into meaningful fragments, which are subsequently interpreted with regards to recurring structural motifs. The analysis relies on the chemical constitution of the molecule (e.g. the protein’s primary sequence of amino acids) and the chemical and crystallographic constraints on the domain (e.g. constraints on how the backbone of the molecule folds in three dimensions).

Although their goals are similar, molecular scene analysis differs from approaches used in machine vision in a number of ways. First, the image for a crystal is perceived and depicted in three dimensions. This eliminates many of the problems of segmentation and recognition involved in vision applications: features in electron density maps are not occluded and we can utilize three-dimensional segmentation and pattern matching techniques. As well, we are not concerned with factors such as light source, surface material or atmospheric conditions, which may distort the

appearance of a visual image. A distinct advantage in molecular scene analysis is the availability of previously determined structures stored in a Protein Data Bank (Bernstein *et al.* 1977). This database of three-dimensional structures forms the basis of a comprehensive knowledge base for template building and pattern recognition in molecular scene analysis; although the scenes we wish to analyze are novel, their parts have almost certainly appeared previously in similar contexts/configurations. Despite overcoming many of the obstacles involved in visual scene analysis, there still exist difficulties in carrying out molecular scene analysis. The primary computational hurdle that exists is the segmentation and interpretation of protein images and is related to the incompleteness of experimental data due to the *phase problem*: image reconstruction is difficult to achieve since phase information is not readily available from experimental data.

In our proposed approach to molecular scene analysis, the process of crystal structure determination is modeled as iteratively resolving the three-dimensional image of the atomic arrangement within the crystal. In this approach, an input image (electron density map) can be analyzed to interpret parts at varying levels of resolution. At low resolution we attempt to identify the molecular envelope; the map is segmented into parts corresponding to residues at medium resolution and the result is analyzed in terms of secondary structure conformations. Only at high resolution can atoms be identified. Any partial interpretation of features of an image at a given resolution provides additional phase information, which can subsequently be used to resolve the current map and provide more information for further analysis.

In summary, the process of determining the tertiary structure of a protein can be likened to a scene analysis, which draws both from the long-term memory of structural motifs and from the application of chemical and

crystallographic rules. In the analysis, the molecular scene is reconstructed and interpreted in an iterative procedure which proceeds from an initially low resolution uninterpreted image to a fully interpreted high resolution map. To accomplish the goals of molecular scene analysis, however, requires the representation of protein structures in a knowledge base that can be easily accessed to retrieve general and specific properties of protein structure at different levels of abstraction (amino acid, secondary structure, molecule, etc.). It also benefits from representing protein motifs in forms that can be visually and spatially compared and interpreted. In the remainder of the paper, we describe a knowledge representation framework that provides the flexibility required for the automated reconstruction of protein structures from crystallographic data.

Representations for Molecular Imagery

To facilitate a computational approach to molecular scene analysis, we incorporate a multimodal framework that has been developed for computational imagery (Glasgow & Papadias 1992; Glasgow 1993b; 1993a). This framework consists of representations and operations that can be used to depict, transform, scan, pattern match and reason with the descriptive and the visual/spatial information contained in images. Computational imagery involves three distinct representations for reasoning with images, each appropriate for a different kind of processing. The *descriptive representation* stores all of the relevant information about a particular domain. The *spatial representation* symbolically models the image in terms of its parts and their relative locations, whereas a *visual representation* for imagery preserves information such as shape, size and relative distance and is generally uninterpreted. Although the functionality of computational imagery was inspired by results in cognitive psychology and neuropsychology, its implementation has focused on expressive power, inferential adequacy and efficiency.

Descriptive representation of proteins: The organization of the comprehensive information for existing protein structures into a knowledge base of molecular scenes forms the cornerstone of our overall strategy for molecular scene analysis. The proposed knowledge base differs fundamentally from the existing crystallographic databases in several respects. First, it contains knowledge, not only in the form of simple facts, but also in the form of general concepts synthesized from the analysis of these facts. Second, it allows the explicit expression of links between knowledge structures through which implicit knowledge can be derived. Third, while crystallographic databases store data in rigid formats organized primarily for search-

ing, the molecular knowledge base is organized to allow for dynamic creation, deletion and modification of knowledge structures and links. Finally, the knowledge base incorporates procedural knowledge to embody the "know how" of experts and existing algorithms.

For many forms of reasoning, it is useful to organize the available information in terms of concepts and objects, with links denoting their meaningful relationships. An obvious and natural organization for molecules is one that reflects their structural hierarchy and shows clearly how molecular scenes can be constructed from (or decomposed into) their building blocks. The network in Figure 2 illustrates several levels of the building block structural (PART.OF) hierarchy of proteins.

At each level of the molecular structural hierarchy, there is another important hierarchy, one that separates the general concepts of the building blocks from the actual instances of these concepts. Historically, the concepts of intermediate structures, such as amino acids, secondary structures, molecules, etc., have developed into general classifications, or categories of structures. These classifications are based on shared features and/or structure. For example, the class of amino acid has several subclasses – one of which is the class valine, as illustrated in Figure 2. This specific class of amino acid has some generic properties (e.g. molecular weight, hydrophobicity, etc.) that are shared by all valines. A particular instance of valine may have unique features (e.g. position in the polypeptide chain, membership in a secondary structure element, etc.) in addition to those it shares with all valines. Thus, we can represent molecules in terms of a second, conceptual hierarchy, based on the classes, subclasses and instances of the building blocks at each level of the structural hierarchy.

The semantic network for proteins has been implemented using a frame system. As well as allowing for the storage of attributes for each class and subclass, this representation also permits the expression of procedures for generating the visual and spatial depictions for both the prototypical structures, and for the particular instances of protein structures or substructures stored in the Protein Data Bank.

Spatial representation of proteins: The model-based approach we consider for representing spatial motifs for protein structures is based on the notion of description logics. Similar to frame-based representations, description logics make a clear division between concepts and instances. Primitive concepts are used in the logic to define more complex concepts by applying constructors such as conjunction, negation and disjunction. For example, assume that *amino-acid*, *hy-*

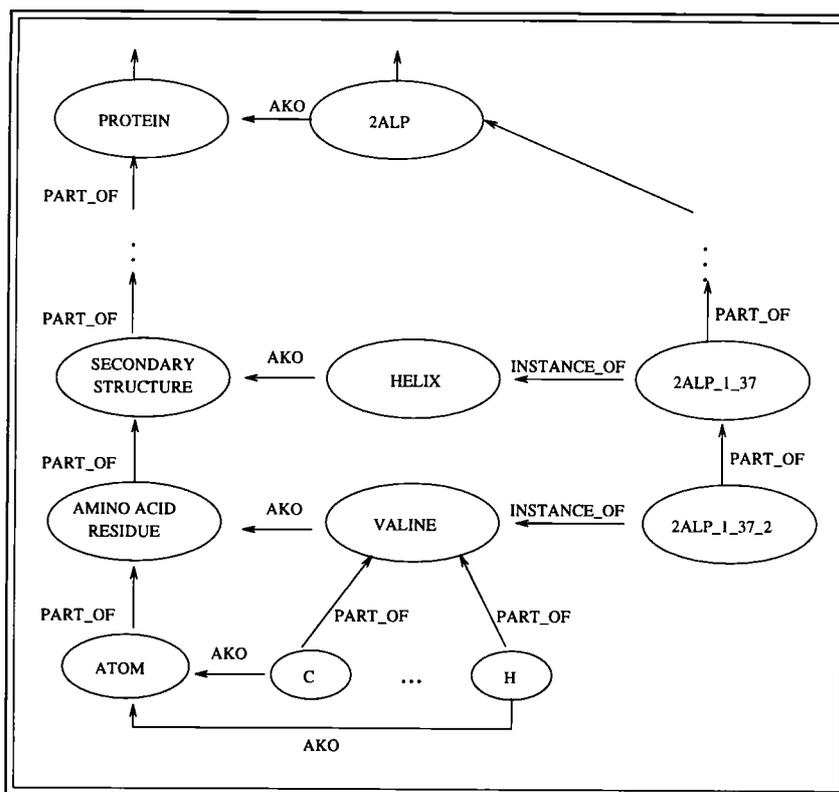


Figure 2: Hierarchical structure of the protein knowledge base

drophobic and *small* are primitive concepts in a logic then the class of amino acids that are small or not hydrophobic could be defined as the concept: (*amino-acid and (small or not hydrophobic)*).

Traditional description logics have limited facilities for expressing relationships between parts of a concept or an instance. Thus, in order to reason about structured objects, we have crafted a *Spatial Description Logic*, *SDL*, which was specifically tailored for the representation, discovery and classification of structured objects (Conklin, Fortier, & Glasgow 1993). *SDL* extends traditional description logics by adding the notion of a *symbolic image*. Such an image is described by a spatial data model comprising a set of concept terms with their coordinates in multi-dimensional space. *SDL* has been used for the representation of *protein motifs*, which are abstractions of observed patterns of amino acids (Conklin, Fortier, & Glasgow 1993; Conklin *et al.* 1996). Figure 3 illustrates a protein structure motif and its symbolic image description in *SDL*. The motif consists of six amino acids, each of which we assume has been specified as a primitive concept. Also associated with a symbolic image is a set of preserved qualitative relations, such as relative distance and spatial configuration.

Visual representation for proteins: Unlike the spatial representation, which is concerned with the relative location of parts of an image, a visual representation is concerned more with quantitative information such as shape, size, angles and distance. Shape information can be extracted from a protein image by contouring the electron density map around density values that exceed a given threshold value. The information in this volumetric representation, however, is often too detailed for efficient visual scene analysis. What is needed instead is a representation that captures the relevant shape information of molecular structures, while discarding unnecessary and distracting details.

A topological approach to representing the visual qualities of a molecular scene was recognized as the most natural way to preserve shape information by capturing the fluctuations of the electron density function of three-dimensional molecular images. Critical points are calculated from the three-dimensional electron density electron map by determining points where the gradient of the density function vanishes. At such points, maxima and minima are defined by computing second derivatives, which adopt negative or positive values respectively.

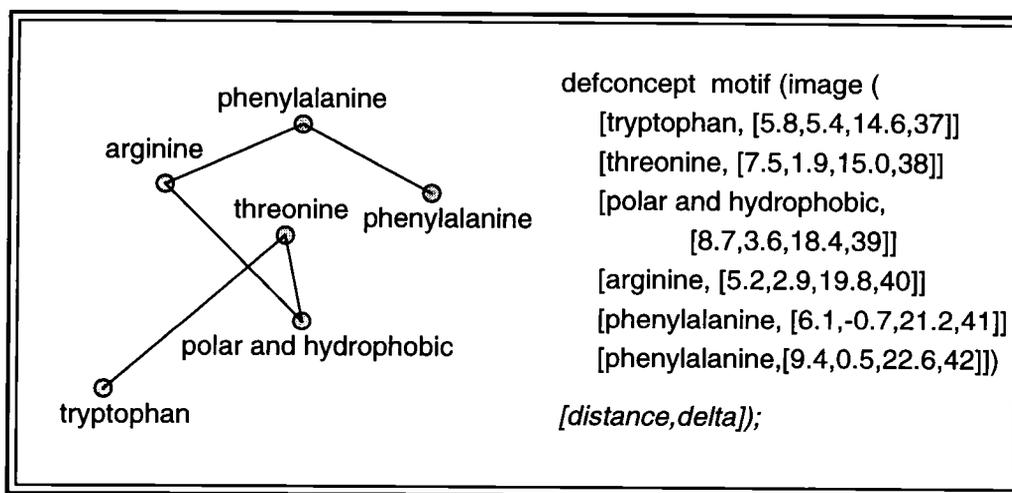


Figure 3: Protein motif described as a symbolic image in *SDC*

We have applied the topological approach to the representation and analysis of medium resolution (3 Å) electron density maps of proteins (Leherte *et al.* 1997). Our studies show that the critical point networks provide a useful segmentation of the electron density maps, tracing a skeleton of the protein main chain of a protein and capturing its overall three-dimensional conformation as a 3D uninterpreted graph (see Figure 4).

The visual representation is primarily used in the low-level analysis of a protein structure. Once we have segmented an image into its critical points and constructed a three-dimensional model of the protein backbone we can further analyse the representation in order to interpret secondary structures (Leherte *et al.* 1994) and identify individual amino acid residues (Baxter *et al.* 1996).

Discussion

The representations for molecular imagery provide a framework for multimodal reasoning in which we can integrate existing sources of protein information, including rules and heuristics used by experts in chemistry, biochemistry and crystallography, the experience accumulated in crystallographic databases, and the methodologies embodied by existing computer programs. This ever-growing knowledge base is organized according to the natural structural and conceptual hierarchies of molecules. The representation of protein knowledge as a semantic network memory model permits rapid retrieval and reasoning with information at all levels of the decomposition hierarchy.

At the core of our knowledge-based approach to molecular scene analysis is the concept of imagery, which involves the ability to represent and reason with

both interpreted and uninterpreted three-dimensional models of molecular structure. The proposed framework for computational imagery was designed to bring to the forefront the fundamental visual and spatial characteristics of a scene. This framework is applied to the domain of protein structures where algorithms have been developed to construct simplified visual representations that depict the shape of an image as a three-dimensional spanning tree. A spatial representation has been proposed that captures and reasons with relationships such as bonding, spatial configuration and local symmetry. Techniques are currently being developed for organizing, retrieving, and pattern matching molecular scenes based on their visual and spatial qualities.

Artificial intelligence research is concerned with the discovery of computational tools for solving hard problems that rely on the use of extensive knowledge. To solve problems in the molecular domain requires ways of representing and reasoning with the varied and vast amount of available knowledge. There is clearly a need for organized chemical and structural knowledge in areas such as macromolecular structure determination and rational drug design. These applications also benefit from the ability to represent, discover, pattern-match and reason with visual and spatial motifs of molecular images.

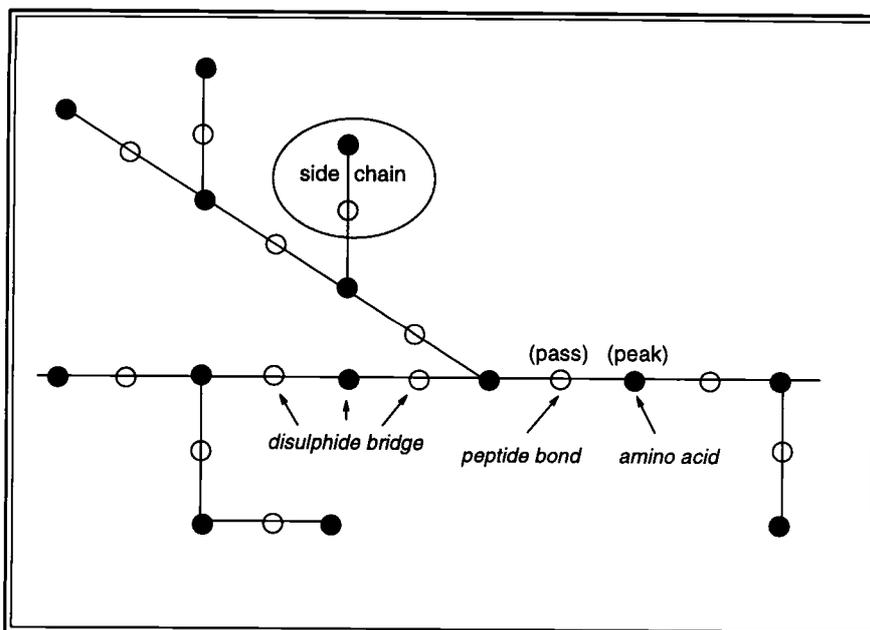


Figure 4: Planar representation of the critical point spanning tree for a portion of the protein structure IRNT.

Acknowledgements

The research presented in this paper was supported by the Natural Science and Engineering Research Council of Canada.

References

- Baxter, K.; Steeg, E.; Lathrop, R.; Glasgow, J.; and Fortier, S. 1996. From electron density and sequence to structure: Integrating protein image analysis and threading for structure determination. In *Proceedings of the 4th International Conference on Intelligent Systems for Molecular Biology*, 25–33. AAAI/MIT Press, Menlo Park, California.
- Bernstein, F.; Koetzle, T.; Williams, J., Jr., E. M.; Brice, M.; Rodgers, J.; Kennard, O.; Shimanouchi, T.; and Tasumi, M. 1977. The Protein Data Bank: A computer-based archival file for macromolecular structures. *Journal of Molecular Biology* 112:535–542.
- Conklin, D.; Fortier, S.; Glasgow, J.; and Allen, F. 1996. Conformational analysis from crystallographic data using conceptual clustering. *Acta Crystallographica B* 52:535–549.
- Conklin, D.; Fortier, S.; and Glasgow, J. 1993. Representation for discovery of protein motifs. In Hunter, L.; Searls, D.; and Shavlik, J., eds., *Proceedings of the First International Conference on Intelligent Systems for Molecular Biology*. AAAI/MIT Press, Menlo Park, California.
- Fortier, S.; Castleden, I.; Glasgow, J.; Conklin, D.; Walmesley, C.; Leherte, L.; and Allen, F. 1993. Molecular scene

analysis: The integration of direct methods and artificial intelligence strategies for solving protein crystal structures. *Acta Crystallographica D* 49:168–178.

Glasgow, J., and Papadias, D. 1992. Computational imagery. *Cognitive Science* 16(3):355–394.

Glasgow, J.; Fortier, S.; and Allen, F. 1993. Molecular scene analysis: crystal structure determination through imagery. In Hunter, L., ed., *Artificial Intelligence and Molecular Biology*. AAAI Press, Menlo Park, California. 433–458.

Glasgow, J. 1993a. Imagery and AI – where do we go from here? *Computational Intelligence* 9(4):424–435. Response to taking issue forum.

Glasgow, J. 1993b. The imagery debate revisited: A computational perspective. *Computational Intelligence* 9(4):309–333. Taking issue paper.

Leherte, L.; Fortier, S.; Glasgow, J.; and Allen, F. 1994. Molecular scene analysis: A topological approach to the automated interpretation of protein electron density maps. *Acta Crystallographica D* 50:155–166.

Leherte, L.; Glasgow, J.; Baxter, K.; Steeg, E.; and Fortier, S. 1997. Analysis of three-dimensional protein images. *Journal of Artificial Intelligence Research (JAIR)* 125–159.

Marr, D. 1982. *Vision*. W.H. Freeman and Company: San Francisco.