ICS-FORTH **1997**

I2Cnet Medical Image Annotation Service

[Short title: A Framework for Providing Medical Services over the World Wide Web]

Catherine E. Chronaki[1], Xenophon Zabulis[1,2],
and Stelios C. Orphanoudakis[1,2]

[1]*Institute of Computer Science, Foundation for Research and Technology - Hellas, P.O. Box 1385, 711 10 Heraklion, Greece*

[2]*Department of Computer Science, University of Crete, Heraklion, Greece*

Tel.: +30 (81) 391419, Fax.: +30 (81) 391601[1]
E-mail: {chronaki,zabulis,orphanou}@ics.forth.gr

URL: http://www.ics.forth.gr/ICS/acti/cmi_hta/publications/papers.html

Medical Informatics, Special Issue,
vol. 22(4), 337-347, 1997

Abstract. I^{2}Cnet (Image Indexing by Content network) aims to provide services related to the content-based management of images in healthcare over the World-Wide Web. Each I^{2}Cnet server maintains an autonomous repository of medical images and related information. The annotation service of I^{2}Cnet allows specialists to interact with the contents of the repository, adding comments or illustrations to medical images of interest. I^{2}Cnet annotations may be communicated to other users via e-mail or posted to I^{2}Cnet for inclusion in its local repositories. This paper discusses the annotation service of I^{2}Cnet and argues that such services pave the way towards the evolution of active digital medical image libraries.

Keywords: web services, medical digital libraries, image annotations

I2Cnet Medical Image Annotation Service

[Short title: A Framework for Providing Medical Services over the World Wide Web]

Catherine E. Chronaki[1], Xenophon Zabulis[1,2],
and Stelios C. Orphanoudakis[1,2]

[1]*Institute of Computer Science, Foundation for Research and Technology - Hellas, P.O. Box 1385, 711 10 Heraklion, Greece*

[2]*Department of Computer Science, University of Crete, Heraklion, Greece*

Tel.: +30 (81) 391419, Fax.: +30 (81) 391601[1]
E-mail: {chronaki,zabulis,orphanou}@ics.forth.gr

Medical Informatics, Special Issue,
vol. 22(4), 337-347, 1997

Abstract. I^{2}Cnet (Image Indexing by Content network) aims to provide services related to the content-based management of images in healthcare over the World-Wide Web. Each I^{2}Cnet server maintains an autonomous repository of medical images and related information. The annotation service of I^{2}Cnet allows specialists to interact with the contents of the repository, adding comments or illustrations to medical images of interest. I^{2}Cnet annotations may be communicated to other users via e-mail or posted to I^{2}Cnet for inclusion in its local repositories. This paper discusses the annotation service of I^{2}Cnet and argues that such services pave the way towards the evolution of active digital medical image libraries.

Keywords: web services, medical digital libraries, image annotations

1. Introduction

The World-Wide Web (WWW) is currently enjoying rapid growth and increasing popularity due to its visual nature and information retrieval capabilities. This has resulted in substantial research effort being devoted to the development and provision of information services. *I^{2}Cnet* (Image Indexing by Content network) is part of a research effort for the development of robust content-based retrieval methods for different classes of medical images.

I^{2}Cnet [1] is organized as a network of *I^{2}Cnet* servers, which interoperate with Image Management and Communication Systems (IMACS) and offer services, related to the content-based management of images. The main requirements of the *I^{2}Cnet* architecture are: *availability* to authorized users throughout the world, *interoperability* with the applications and services in the evolving regional healthcare network of Crete, and *network transparency* that allows users to request a service by name without being aware of the underlying server infrastructure. *I^{2}Cnet* provides a variety of services ranging from intelligent image database browsing and image processing to collaboration facilities.

The image annotation service of *I^{2}Cnet* provides the means for creative interaction with the images stored in the repositories of *I^{2}Cnet* servers and active collaboration with the users accessing and maintaining the repositories. The objective of the annotation service is to provide healthcare professionals with the ability not only to access medical image collections, but also to interact with images, authoring, viewing, and communicating annotations on groups of images. Specifically, the medical image annotation service of *I^{2}Cnet* allows specialists to create annotations on groups of images, communicate them to any Internet or intranet user by e-mail, and search for medical images and/or annotations based on annotation text and image content. Annotation objects created using the annotation service of *I^{2}Cnet* may serve as the basis for collaboration with other specialists through the discussion and exchange of ideas and opinions. The annotation service of *I^{2}Cnet* is accessible through any Java-enabled browser. In this way, authorized users around the world may access the services of *I^{2}Cnet* and interact with the material present in the *I^{2}Cnet* repositories.

The rest of the paper is organized as follows: Section 2 further elaborates on the concept of annotations as they are employed in *I^{2}Cnet*, establishing a case for annotations in a world of evolving digital medical libraries. Section 3 discusses the issues involved in designing an

annotation system and introduces the annotation services of *I²Cnet* focusing on design, architecture, functionality, and user interface issues. Section 4 addresses collaboration aspects of the annotation service. Section 5 deals with the problem of interoperability as it is addressed by *I²Cnet* services. Section 6 overviews related work on annotation systems and related web services. Finally, section 7 concludes the paper with our current and future work.

2. Annotations: A way to talk about sets of images

An Internet connection provides healthcare professionals at remote locations with the necessary infrastructure to reach the regional healthcare network, get the latest updates in their specialty, read news, contact their peers, access on-line medical journals, or browse medical image collections. *I²Cnet* services also allow them to browse through medical image archives, guided by different notions of image similarity, and interact with the contents of these archives. In particular, the medical annotation service of *I²Cnet* allows users to interact with information that is mostly static, associating it with dynamically created entities, i.e. annotations.

Frequently, a healthcare professional that browses through an image collection wishes to select a few images and create a note, possibly private, pointing out interesting features and comparing these images to images from other collections. Such annotation objects may serve as the basis for communication with other specialists, discussion, and exchange of ideas and opinions. Thus, the *I²Cnet* annotation service adds a dynamic dimension to the image collection by allowing users to create annotations that relate image sets to text, graphics, and other forms of medical data. Annotation objects that are created in the context of *I²Cnet*, can be communicated in HTML format, to any Internet or intranet user. Furthermore, the annotation service, when used in combination with the image posting service of *I²Cnet*, will allow authorized users to post interesting medical cases as annotations into *I²Cnet* servers. Thus, digital medical image libraries will begin to evolve with contributions by healthcare professionals.

2.1 What is an image annotation?

An image annotation is simply a way to talk about images. Annotations provide people with the means to communicate and exchange ideas about what they see and think. Currently, image annotations are gradually making their appearance in various contexts ranging from educational systems, to digital libraries and collaboration fora. In each context, an annotation may serve a different purpose, facilitating rich forms of interaction with digital objects and collaboration among the people accessing and maintaining the relevant information. Specifically, the purpose of an annotation may be:

- to *comment* on or point out specific attributes or features of the image.
- to *confirm* a comment made by another user, i.e. seals of approval.
- to *refute* a comment made by another user.
- to *illustrate* a scenario, possibly via trails of annotations.
- to *correct* a mistake.

Furthermore, annotations play an increasingly vital role in the digital world because they provide a means of asynchronous collaboration, complementary to e-mail, news, bulletin boards, discussion fora, etc [2].

2.2 Image Annotations in I^{2}Cnet

Most equipment associated with various medical imaging modalities, such as Computed Radiography (CR), Computed Tomography (CT), Magnetic Resonance Imaging (MRI), Ultrasound, etc, provide textual annotations and graphic overlays as the means to comment on the contents of a medical image, such as a tumor, a hemorrhage, a pregnancy, etc. In contrast to these systems that incorporate just textual annotations and graphic overlays, *I^{2}Cnet* is based on a broad notion of annotations. First, *I^{2}Cnet* annotations may apply to a group of annotated images, instead of a single image. In this way, the user may compare multiple images or comment on a trend illustrated by a sequence of images. Another aspect of *I^{2}Cnet* annotations is that each annotation is associated with *pointers* indicating relevant digital objects. Pointers are uniform resource locators (URLs) which may point to a multimedia document, an audio file, or another annotation.

As shown in Fig. 1, an *I^{2}Cnet* annotation includes one or more *details* which are graphical objects specifying a region of interest in an annotated image. Frequently, a detail is linked to *notation text*. Notation text is a textual field, which may refer to the region of interest outlined by the detail or the annotated image as a whole. Multiple details and notation texts may be grouped together to form *overlays*. Overlays may be viewed concurrently or one at a time. Additionally, an *annotation message* enables the inclusion of a text report or message that supplements the graphical part and the *pointers* of the annotation. This message may include an overview or the purpose of the annotation. Finally, each annotation is associated with a *property* object, which maintains meta-information including the author of the annotation, creation date, review date, moderator, subject, and relevant keywords. Meta-information is used for annotation indexing and search.

This broad view of annotations facilitates the use of the annotation service in a variety of application contexts. One such application concerns the authoring and maintenance of medical case reports. A medical specialist may use the annotation services of *I^{2}Cnet* to create his private library of medical cases. For example, an annotation may illustrate the evolution of a brain tumor in the course of therapy. The list of pointers may associate auxiliary lab reports or comments made by the physician at various stages. Other pointers may indicate similar cases that have appeared in medical journals.

Another noteworthy application is pertinent to educational systems. A teacher may use *I^{2}Cnet* annotations to develop questionnaires regarding the anatomical structures present in the annotated medical images. The annotation overlays enable the gradual introduction of additional features on the image, while pointers refer to additional questionnaires or monitored resources. The annotation message provides general information on the questionnaire. Thus, *I^{2}Cnet* annotations allow user interaction and collaboration, while they provide the means for visual information organization.

3. Annotation Services in *I^{2}Cnet*

3.1 Design Issues

This on-going research activity has raised several design issues concerning annotation environments associated with repositories of digital information.

Annotation types. In general, annotations may be *public*, *shared*, or *private*. Currently, *I²Cnet* supports all three types of annotations. Public annotations are those posted to *I²Cnet* and linked to the image repository of an *I²Cnet* site. Private annotations are those authored by a user using the annotation service of *I²Cnet* and maintained in private disk space. Finally, annotations may be implicitly shared if a user chooses to explicitly communicate an annotation to one of his/her peers, or maintains an annotation in shared network space (virtual shared workspace).

User Roles. An annotation environment may accommodate a variety of user roles according to the problems it addresses. Typical user roles in an annotation environment are *reader*, *author*, and *moderator*. A specific user may take a different role with regard to a different repository or annotation. In *I²Cnet*, for each site there is a moderator who reviews the public annotation material posted to that site. When an annotation is posted for review, the *moderator* may choose to publish the annotation or discard it. The only applicable role concerning the public annotations linked to images in an *I²Cnet* repository, is reader. Even the author of an annotation may not modify the annotation once it is posted to *I²Cnet*.

Supported Operations. Basic operations supported by annotation environments are authoring, storage, search, and communication. Each annotation is associated with meta-information (attributes), which facilitate indexing and keyword search. More advanced global annotation systems also support annotation expiration and semantic search based on the contents of the annotation [12]. In its initial version, *I²Cnet* supports only the basic set of operations. As the system evolves, these basic operations will be refined and additional ones, such as semantic search, will be introduced.

Presentation. Various systems indicate the presence of annotations in different ways. In general, there are *in-line* annotations, which are indicated with special marks on the annotated object [5], or *out-of-line* annotations whose presence is indicated by special marks at specific points or next to the annotated object [2]. Presently, the HyperNews type of presentation [9] enjoys significant popularity. The HyperNews style of presentation is reminiscent of threaded news. Next to the annotated objects, a multilevel hyperlist provides access to related annotations. In *I²Cnet*, users browsing an image class in an *I²Cnet* repository view image miniatures. A special icon next to an image miniature indicates the presence of annotations. By clicking on that icon, the list of public annotations associated with the specific image appears.

Discourse Structures. Discourse structures refer to the associations made among annotations in various annotation systems. Again, discourse structures analogous to those supported by HyperNews seem to be gaining wide acceptance. *I²Cnet* takes an unsophisticated approach regarding discourse structures: it provides the basic mechanisms for associating annotations, even though it does not advocate specific policies. Annotations may be linked to each other through the use of pointers. As the annotation service becomes more widely used, we expect that specialized discourse structures will be introduced.

Lifetime of annotations. Some annotation systems associate annotations with a special attribute called *expiration time* [9]. When the annotation is posted, the suggested expiration date indicates the lifetime of the annotation as proposed by its author. When annotations are part of a repository, the administrator responsible for the maintenance and the content of the repository may purge annotations at his discretion. Currently, there is no *expiration time*

attribute in *I²Cnet*. The repository administration is responsible for the annotations published in the repository.

User notification. Some users might like to be notified when a new annotation is linked to a specific image. Some annotation environments such as Commentor [5,8] and Grassroots [10] provide this service. In *I²Cnet*, there is currently no provision for services of this type.

Scalability. An annotation system is scalable if it can accommodate a large number of annotations and a large number of user requests. Researchers involved in the Mosaic development team [4] have addressed the scalability problem for annotation systems. The limiting factor, regarding the scalability of the *I²Cnet* annotation service, is the scalability of the HTTP protocol and the size of the annotations authored by the user. The *I²Cnet* annotation format is relatively compact due to its descriptive and associative nature, resulting in small annotation objects.

3.2 I2Cnet Annotation Server

Each *I²Cnet* server maintains thumbnails, descriptions, and annotations of medical images. Images are organized into high level classes that express medical imaging modalities and parts of the anatomy. By selecting an image class, the user may browse through thumbnails of the images that belong to that class. Next to each thumbnail, a special icon indicates the presence of public annotations. When the user clicks on that icon, a hyperlist of associated annotations appears. Each item on the hyperlist specifies meta-information on the annotation such as posting date, subject, author, etc. When an item is selected, the annotation service is requested to enable access to the selected annotation. The user request is directed to the web server of the *I²Cnet* daemon and the annotation applet is loaded in the browser of the user. At initialization time, the annotation applet establishes a dedicated connection to the *I²Cnet* daemon and, from that point on, the user may author, save, convert, and communicate annotations.

A notable aspect of *I²Cnet* services in general, and the annotation service in particular, is that they take advantage of proxy and cache services configured for particular users. This means that the annotation applet, and generally all data used by the applet, may be present in the cache of the browser, or a proxy in the proxy hierarchy between the user and the host *I²Cnet* server, thus conserving bandwidth and latency (Fig. 2).

3.3 User Interface

The user interface of the annotation environment consists of a graphical area, toolbars, a text window, and the current workspace (Fig. 3). An *I²Cnet* annotation on a group of images consists of images to be annotated, graphical objects (drawings and text), an annotation message, and pointers (audio, video, related annotations, etc). The user may request information on the components of the annotation on display: images, text, and drawings. Information on drawings includes geometric properties, while information on images includes owner, size, origin, type, etc.

When creating an annotation from scratch, a set of one or more images are selected from the workspace for addition to the graphical area. Then, using the set of available drawing tools, the user may add geometrical shapes (lines, rectangles, ovals, etc.), hand-drawn shapes, arrows, and text. Individual shapes may be grouped and treated as a compound object. Each

object can be edited, moved, or deleted while its properties may be previewed. In addition to text that may be typed on the graphical area, a separate text window provides space for detailed comments to be included in the annotation message. The resulting annotation object may be posted to the *I²Cnet* server, exported in HTML format, or mailed to other Internet users by e-mail.

3.4 Storage Management

Each *I²Cnet* annotation is stored in ASCII format to enable content-based search. The introduction of other storage formats, e.g. CAD DXF [3] and laby [11] that enable the portability of *I²Cnet* annotations across different systems, has also been considered. Annotation files include references to the annotated images, URLs, graphical objects (text, lines, drawing, etc.), the free text of the annotation message, and the list of pointers. In addition, for each annotation, *I²Cnet* maintains a separate *property* file, which stores the meta-information of the annotation. The separation of the annotation material in component files enables the deployment of different types of search.

The *property* object of an image stored in an *I²Cnet* repository includes the list of all annotations associated with it. Every time a new annotation is published in an *I²Cnet* server, the properties of all annotated images are updated. Thus, a user that requests the list of annotations associated with a specific image always gets up-to-date information.

Since all public *I²Cnet* annotations can be modified only by the moderator(s) of the *I²Cnet* site, no database engine is involved in accessing *I²Cnet* annotations. The web browser of the *I²Cnet* annotation server collects statistics of access and use. Concurrency control problems are avoided by ensuring that, at any point, just one user i.e. the site administrator, may be updating the material of the specific site.

3.5 Searching for Annotations

The textual contents of the annotation files are indexed in order to enable various types of search. A user may search on the contents of the annotation message, the graphical text of the annotation, and the meta-information associated with the annotation.

4. Collaboration Facilities

The annotation service of *I²Cnet* is an important means towards information organization and collaboration among healthcare professionals. Annotation objects, created in the *I²Cnet* environment, may be converted into HTML and e-mailed in read-only format to any Internet or intranet user. Alternatively, annotation objects may be e-mailed as *I²Cnet* annotations and be viewed in the context of the *I²Cnet* annotation environment.

Synchronous collaboration is also possible through the on-line collaboration service of *I²Cnet.* In this case, multiple users may use an interactive talk facility, while they concurrently view a white-board which shows a specific annotation. They may draw on the white-board and specific snapshots of their activity can be recorded.

Furthermore, a group of users distributed over the world may collaborate based on the material included in a shared workspace. Shared workspaces support the cooperation of authorized users by providing secure and transparent access to a heterogeneous data collection. This collection includes service results and multimedia data objects (images, voice, video, patient annotations, etc) that were inserted by the users sharing the workspace.

5. Service interoperability in *I²Cnet*

5.1 I²Cnet Services

I²Cnet provides integrated services, such as image processing, content- and annotation-based search for images and annotations, authoring of annotations and image descriptions, and user collaboration, to authorized users through an unmodified web browser. Services are integrated in the sense that both asynchronous collaboration through e-mail, postings, and shared workspaces, as well as synchronous collaboration through an on-line collaboration service, are facilitated.

I²Cnet servers deliver *I²Cnet* services over the WWW. An *I²Cnet* server consists of a Web server that distributes service applets and information content, and an *I²Cnet* daemon that communicates with the service applets. The communication of an *I²Cnet* service with a particular *I²Cnet* server employs well-defined protocols, which ensure the transfer of minimal amounts of information. The resulting bandwidth-conserving architecture allows users to interact with certain parts of the system, unaware of how other services operate. Thus, the learning cycle for potential users is short, and users can easily follow the evolution of particular services.

The user interface of an *I²Cnet* server is composed of web pages that contain applets specific to the provided service. A virtual workspace is assigned to each user that visits an *I²Cnet* site and its contents are displayed on every *I²Cnet* page throughout navigation. Even though the objects included in a virtual workspace are, in general, distributed over a number of *I²Cnet* repositories, the user is able to manipulate them as if they were physically present in the workspace. Since the current workspace is mapped onto all user pages during each session, users have all items of current interest readily available. Furthermore, workspaces persist between user sessions, among services, and among users. Thus, virtual workspaces facilitate service interoperability, persistence, and user interaction in *I²Cnet*.

5.2 Virtual Workspaces

In *I²Cnet*, service interoperability is accomplished through virtual workspaces, which provide access to intermediate service results in addition to other forms of multimedia data. In this way, the user is able to retrieve from the workspace the result of a service request and use it as input to another service. For example, the user may apply a segmentation algorithm to an image and use the segmented image as the input to an image description algorithm. The resulting image description may, then, be used as a visual query to a content-based retrieval service (Fig. 4).

Virtual workspaces in *I²Cnet* follow the “network computer” paradigm of information management, securing transparent access to the data the user needs, anytime and anywhere, given that an Internet connection and browser software is available. By logging into a workspace, the user may view the items that were inserted into the workspace during his last interaction with the workspace and preview the property sheet of each data object, which includes meta-information such as its author, the date it was last updated, and a brief description. A workspace maintains an index of Uniform Resource Locators (URLs) for the objects it contains. This fact, in combination with the paperless office metaphor that is used in the user interface of the virtual workspace, gives the element of network transparency to the system and increases its usability.

All service pages in a specific user session have access to the same workspace and, therefore, the seamless transfer of data from one service to another is transparent regardless of the actual location of the data object in *I²Cnet* or the Web. Furthermore, multiple users may collaborate over a workspace, sharing material of common interest. Multiple users may connect to the same workspace, for the purpose of short-term cooperation and be notified of any workspace updates as they occur. In this case, the workspace operates as a virtual blackboard and provides the basic infrastructure for Computer Supported Cooperate Work (CSCW), since, when one user adds items to the workspace, these items are immediately visible to other users in the same workspace.

6. Related Work

A number of researchers from different disciplines [2,4-8] have addressed the problem of maintaining annotations on the web. Reported work concerns mostly document annotations and focuses on GroupWare activities. The NCSA Mosaic browser supports private, user, and workgroup annotations on documents by associating URLs with comments [4]. CoNote [2] was a project at Cornell, which provided support for annotations in class material. Another system developed in the context of the Stanford digital library project is Commentor [5,8]. Both CoNote and Commentor address the issues of document and annotation sharing, security, organizing annotation into conversations, searching annotations by content, and annotations referencing other annotations. A good review of issues relevant to the maintenance of public annotations on the web appears in [7].

Even though our work addresses similar issues, it focuses on image rather than document annotations. In addition, the fact that the annotation service deals with medical images introduces additional quality of service requirements. Another important issue affecting the design of the *I²Cnet* annotation environment is interoperability with other applications and services in the regional healthcare network. In this context, the annotation service serves as a customizable architectural component, providing users with the ability to interact with visual information.

7. Conclusions - Future Work

This paper presents the annotation service of *I²Cnet*, which allows users to create annotations on groups of images and communicate them to other users via e-mail or by publishing them in *I²Cnet*. Our current work focuses on searching annotations by image and textual content and combining visual and textual information associated with annotation objects. We are also considering the introduction of different types of annotations and support for discourse structures.

The integration of the annotation service with the patient record management system of a health telematics regional network and the provision for synchronous communication using the annotation environment, are also part of our current work. In this way, the *I²Cnet* annotation service will not only add a active dimension to user interaction with the information space of the regional healthcare network, but it will also promote the sharing of opinions, knowledge, and expertise.

References

[1] Stelios C. Orphanoudakis, Catherine E. Chronaki, and Despoina Vamvaka. I2Cnet: Content-based Similarity Search in Geographically Distributed Repositories of Medical Images. Journal of Computerized Medical Imaging and Graphics, Vol. 20, No. 4, pp. 193-207, 1996.

[2] Jim Davis and Dan Huttenlocher. Shared Annotations for Cooperative Learning. In Proceedings of Computer Support for Collaborative Learning '95, October 17-20, 1995, Indiana University, Bloomington, IN.

[3] DXF format file specification, Appendix C, AutoCAD manual. URL: http://wwwis.cs.utwente.nl:8080/~faase/DWG/dxf.html.

[4] Daniel LaLiberte and Alan Braverman. A Protocol for Scalable Group and Public Annotations. In Proceedings of the Third International World-Wide Web Conference, April 10-14 1995, Darmstadt Germany.

[5] Roscheisen, C. Mogensen and T. Winograd. Beyond Browsing: Shared Comments, SOAPs, Trails and On-line Communities. In Proceedings of the Third International World-Wide Web Conference, April 10-14 1995, Darmstadt Germany.

[6] Stanford Digital Library Project: URL: http://walrus.stanford.edu/diglib/

[7] Gramlich. Public annotation systems, 1994. URL: http://playground.sun.com:80/~gramlich/1994/annote/.

[8] Matthew A. Schickler, Murray S. Mazer, and Charles Brooks. Pan-Browser Support for Annotations and Other Meta-Information on the World Wide Web. In Proceedings of the Fifth International World Wide Web Conference, May 6-10, 1996, Paris, France.

[9] Daniel LaLiberte: HyperNews Home Page, A prototype of our public annotation system. URL: http://union.ncsa.uiuc.edu/HyperNews/get/hypernews.html

[10] Kenichi Kamiya, Martin Roscheisen, and Terry Winograd. Grassroots: A System Providing A Uniform Framework for Communicating, Structuring, Sharing Information, and Organizing People. In Proceedings of the Fifth International World Wide Web Conference, May 6-10, 1996, Paris, France

[11] Christos Georgis. Design and Implementation of the Graphics Portion of the Labyrinth System. Technical report no. 51, M.Sc. Thesis, University of Crete, Greece.

[12] Panos Constantopoulos and Martin Doerr. The Semantic Indexing System. Internal Report, Institute of Computer Science, Foundation for Research & Technology Hellas; 1995.

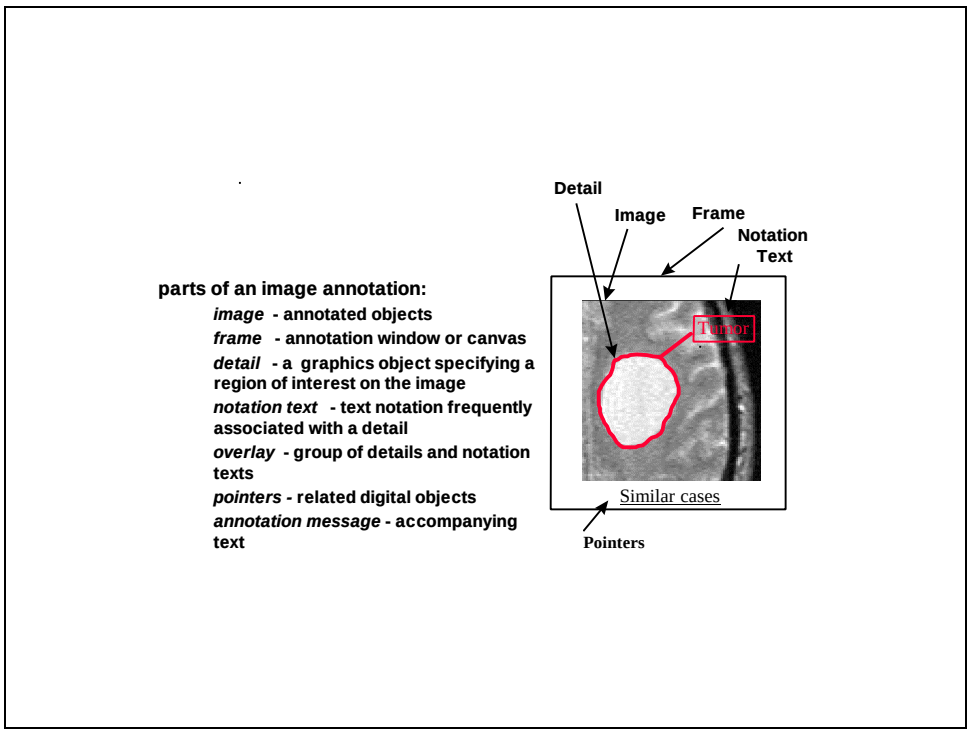

Figure 1: *I²Cnet* image annotation format.

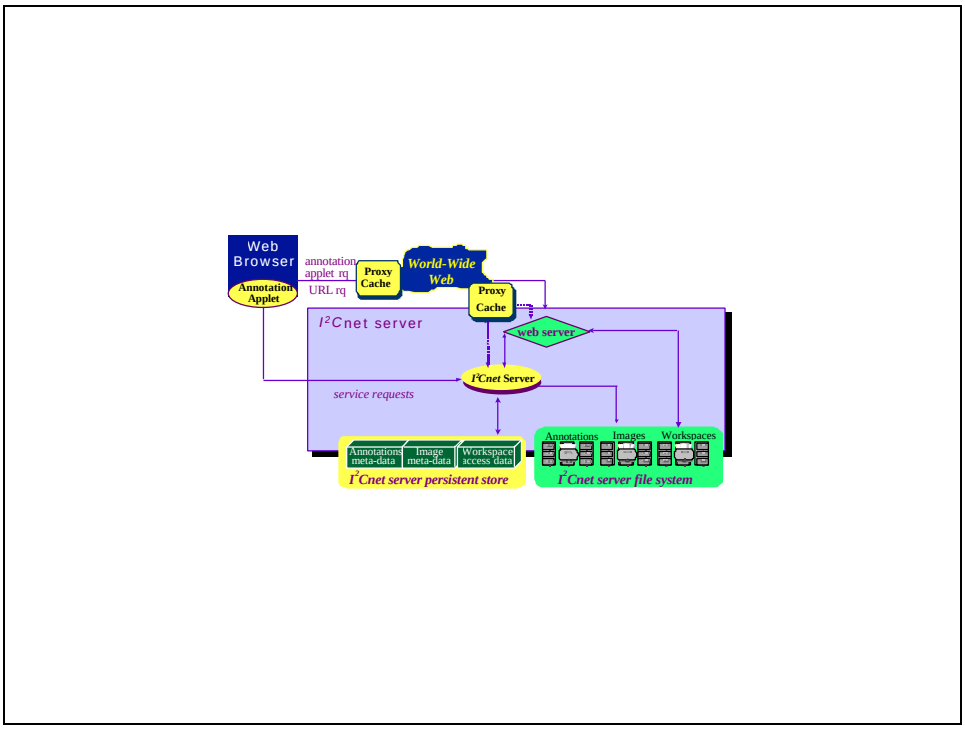

Figure 2: Architecture of an *I²Cnet* annotation server.

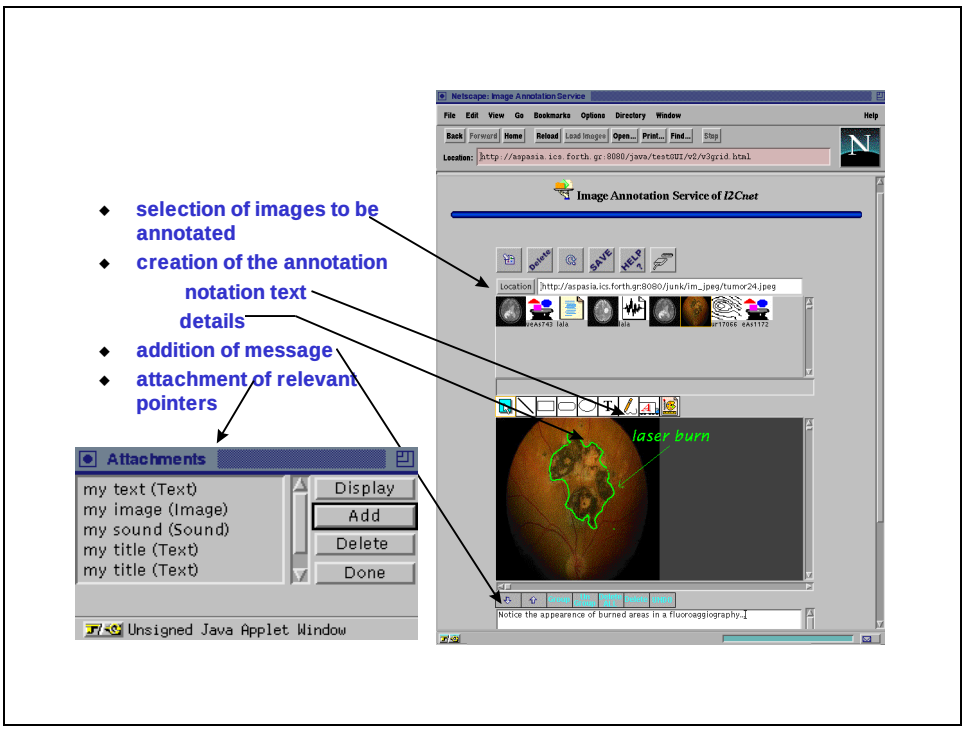

Figure 3: User interface of the *I²Cnet* annotation service.

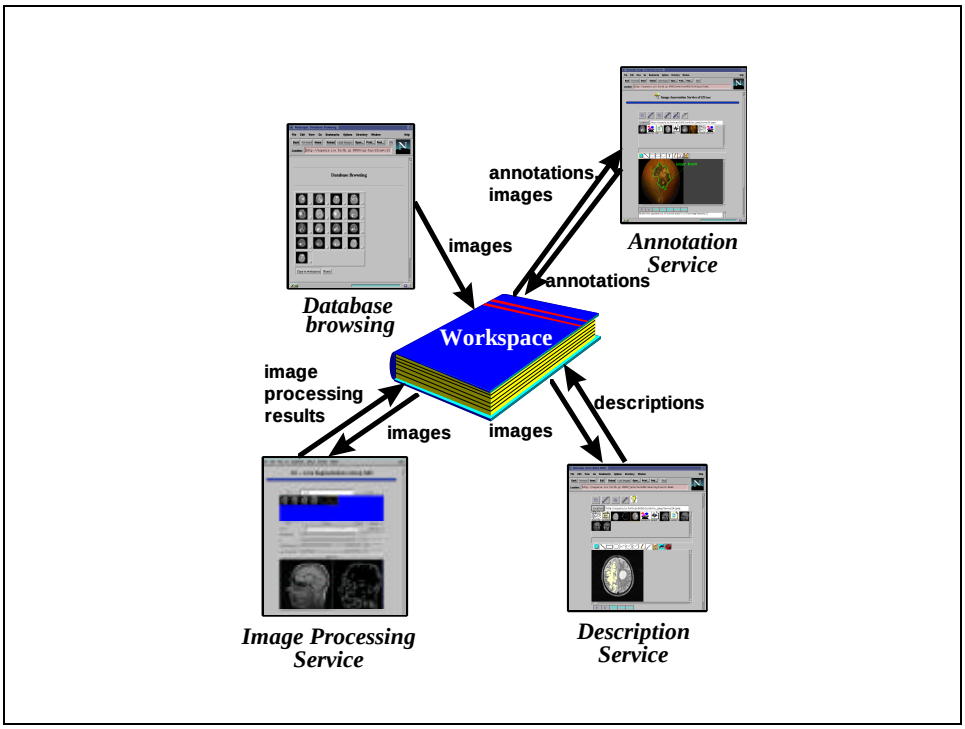

Figure 4: In *I²Cnet*, service interoperability is attained through virtual workspaces.