9

The '*space*' in spatial assistance systems: *Conception, Formalization, and Computation*

MEHUL BHATT, CARL SCHULTZ,
AND CHRISTIAN FREKSA

9.1 The spatial premises: introduction and overview of chapter

Spatial thinking, conceptualization, and the linguistic communication of commonsense as well as expert knowledge about the world is one of the most important aspects of everyday human life. Philosophers, cognitive scientists, linguists, psycholinguists, ontologists, information theorists, computer scientists, and mathematicians have each investigated space through the perspective of the lenses afforded by their respective field of study. Interdisciplinary studies on spatial cognition, for example 'language and space', 'spatial memory', 'spatial conceptualization', 'spatial representations', 'spatial formalizations', 'spatial reasoning' are extensive and enormous to say the least. Within this book itself, other chapters present an elaborate review of the state-of-the-art for some of these fields of study.[1]

We address 'space' from a formal modelling and computational viewpoint, that is, space, as it is interpreted within the computer science disciplines concerned with the investigation of artificial intelligence and knowledge representation in general, and formal methods in spatial representation and reasoning in particular (Freska, 1991; Aiello et al., 2007; Cohn and Renz, 2007; Bhatt et al., 2011). Indeed, the connections

* We gratefully acknowledge the funding and support of the German Research Foundation (DFG), www.dfg.de/. The work described in this paper has been conducted as part of the DFG funded SFB/TR 8 Spatial Cognition project [DesignSpace], www.sfbtr8.spatial-cognition.de/designspace.html.

[1] In particular, chapters most directly related our work include: (1) the chapter by Gallay et al. on navigation assistance systems for blind pedestrians; (2) the chapter by Barclay and Galton on reference object selection for locative expressions; and (3) the chapter by Taylor and Brunyé on the cognitive underpinnings of survey and route description comprehension. Further comments on the relationship with these works are included as they accrue in the rest of the chapter.

between formal models of space, and space as it occurs in language, are themselves a matter of intense research activity, for example within initiatives concerned within generalized linguistic ontologies of space—this is an aspect that we do not directly address in this chapter. However, the chapter does not lose sight of the inherently 'linguistic' aspects and presents several examples of spatio-linguistic discourse material within a broad range of case studies coming from domains with a real application impact. Against the backdrop of this material, we build practically grounded case studies and strive to concretely present the relationship between the conception, formalization, and the computational aspects of '*space*', as it occurs within systems of human assistance, assurance, and empowerment.

9.1.1 Assistance systems

The core of our perspective in this chapter is rooted in our interpretation of a general class of systems concerned with assistance, assurance, and empowerment of humans in their everyday professional and personal lives. Confining the discussion to the spatial domain, these are systems of human–computer interaction involving the representation of space from several different perspectives—the psycho-linguistic and cognitive perspective of humans, the abstract knowledge-centric (symbolic) perspective of an intelligent software construct, and the all too complex and inherently numerical or quantitative perspective of the computer. In our terminology, spatial assistance systems are cognitive agents that 'know' the properties of physical space and are skilled to deal with them in such a way that they can support humans. A special requirement for spatial assistance systems is that they are able to empathize with their human partners to a certain extent; that is, they should adapt to the needs of people rather than require people to adapt to their needs.

9.1.2 Multi-perspective semantics

In order for (our select class of) assistance systems to achieve their functional objectives of human assistance, assurance, and empowerment, the perspective of each stakeholder within the human–computer interaction process has to be adequately accounted for. This must be achieved in a manner of representation and ontological conceptualization that is consistent with the respective *spatial* interpretation of a specific stakeholder. We refer to this notion as *multi-perspective representational semantics*, or simply, multi-perspective semantics for short.

Take the case of an *architecture design assistance system*: here, a designer conceptualizes and abstracts the *structural form* of an environment with respect to an abstract and innate understanding of the shape, layout, and the connectivity of a *spatial structure*. For instance, the *designer* represents one perspective consisting of concepts and relationships from the design domain. By a complex mysterious process we term 'creativity', the designer's abstract notion of structural form is then translated into a real design. She or he uses a tool for spatial design in the form of a floor plan. This plan contains a geometric feature hierarchy consisting of points, line segments, and polygons. These geometric elements are a part of what constitutes another perspective with the system, namely, the perspective of the *design tool*. Now imagine a symbolic design reasoning module that is aimed at deriving the independent as well as inter-related inferences about the perspectives of the designer and the design tool. For example, the reasoning system may be entrusted with the responsibility that a particular design requirement such as 'spacious' or 'private' in the context of a particular 'Room' within the design are indeed satisfied. Here, the conceptual constraints expressed by the designer need to be validated with respect to the realizations within the design tool, with respect to the perspective of the *design reasoner*. In the case of this example, the design reasoner utilizes its own perspective, minimally consisting of qualitative abstractions about the topological, relative orientational, and other spatial aspects of the quantitatively modelled design.

What multi-perspective semantics suggests is that the ontological viewpoints of each of the stakeholders—*designer, design reasoner, design tool*—involved within the assistance system are important, and have to be accounted for in their own right. Taken together, these perspectives constitute the essential nature and character of 'space' within a spatial assistance system. Multi-perspective semantics is further elaborated in Section 11.4.

9.1.3 Language as a representation of structural form, and behavioural function

Continuing with the architecture design assistance scenario, consider an architect or an interior designer confronted with the objective to design an office environment in response to a pre-specified set of client requirements. The expert designer, either individually or within a group interaction, conceptualizes the design task bearing in mind the structural form of the environment with respect to the corresponding functional expectation, which the conceptualization is expected to produce. That is, a structure is envisioned with respect to an anticipated behaviour, and the behaviour or a set of behaviours satisfy a desired function.[2]

9.1.3.1 Interior space description As an example, consider either an architect, an interior designer, or more ambitiously, a system concerned with interpreting, annotating, or understanding the spatial semantics of an interior space description:

As you enter the office, there are two tall bookcases directly to your left. Sitting on the top of the bookcase <u>farthest from you</u> is a potted plant that hangs down almost to the floor....In front of you from the door is a small <u>round</u> table with two chairs <u>on either side</u> of it....

[2] This view of the *functional* aspects of a design bears close relationships to the ontological framework of the Function-Behaviour-Structure (FBS) (Gero, 1990; Gero et al., 1991) model of the design process.

Further into the room is a small leather couch. A desk is directly behind the couch. There is also a floor lamp between the couch and the white board. The four bookcases to the right of the couch extend along that wall so that the last one is parallel to the desk. Across from that bookcase, on the other side of the desk, is one last tall bookcase....There is an open path between the couch and the bookcases to get to the far end of the office. Across from the large table on the same wall as the whiteboard is a second desk that is arranged perpendicularly to the first desk.... The back wall of the office is mostly composed of a large window that extends almost to the floor. A radiator with a flat top is under the window. Several more potted plants are sitting on the ledge of the radiator as well as some coffee mugs and a coffee maker. (Volen 258 Center Description, ISO-Space 2010 Workshop Documents (Pustejovsky, 2010).)

We have derived a design of the interior description on the basis of the Volen 258 text (the text was provided as a handout at Airlie ISO-Space workshop (Pustejovsky, 2010)), illustrated in Figure 9.1: Figure 9.1(a) is a two-dimensional floor plan corresponding to the spatial description, whereas Figure 11.1(b) is its corresponding three-dimensional interpretation. Finally, Figure 9.2 illustrates the spatial descriptors—features and relationships—that essentially determine the overall constitution of the interior space. It is easy to see the crucial role of paths and patterns such as 'circularity', 'extends_along', 'further_into', 'open_path', 'composed_of'.[3] From the viewpoint of computing spatial relations for design, the spatio-linguistic, semantic, qualitative, and quantitative interpretation of descriptors such as these acquires a significant role within systems, that is, automated reasoning processes. In the domain of design assistance, these processes are concerned primarily with two key aspects:

- *design engineering*: this refers to the creation of a structural form that conforms to a behavioural and functional specification;
- *design reverse engineering*: given a structural specification, this corresponds to inferring the extent to which a precise structure fulfils a set of anticipated functional requirements. In this form of analysis, often the easier approach is to look for malfunction, as opposed to an extensive enumeration of the functional aspects.

The above described notion of structural forms and their corresponding behavioural and functional entailments is not confined to the domain of architectural design. In this chapter, we will illustrate the utility of this line of thought to a completely different design domain, namely the creative design of media. Furthermore, outside of the design domain, we also illustrate their applicability for the domain of real-time emergency assistance.

9.1.4 Aim and Organization

[3] The design exercise illustrated in Figure 9.1 has been made taking into account the 'narrative-like' structure of the text as a whole, and a consolidation of all the 'perspectives' that are offered therein.

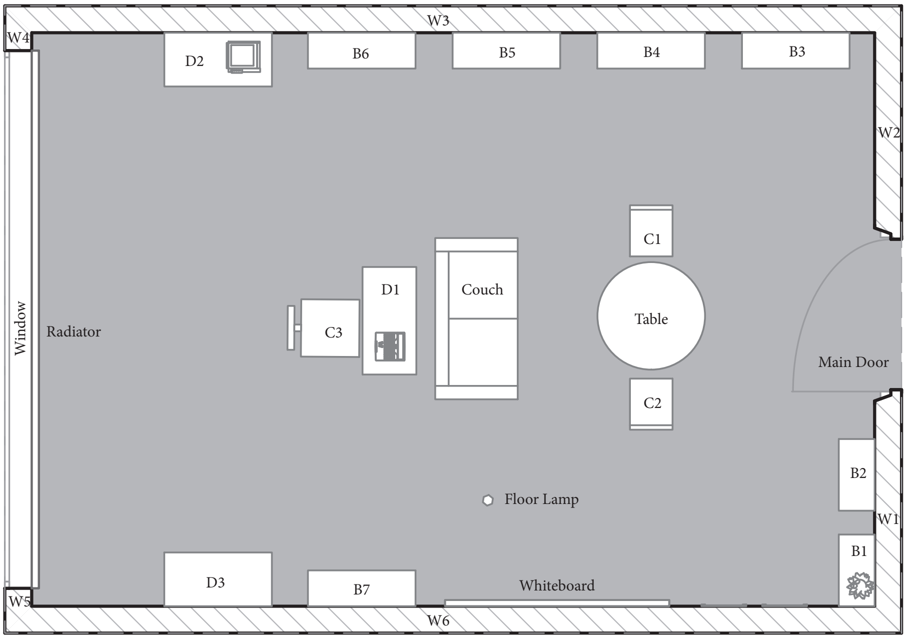

(a) 2D Floor Plan

(b) 3D Conceptualisation (view from the window)

FIGURE 9.1 *Our interpretation of the Volen 258 Center Design*: Role of Paths, Patterns and Commonsense Knowledge in Design

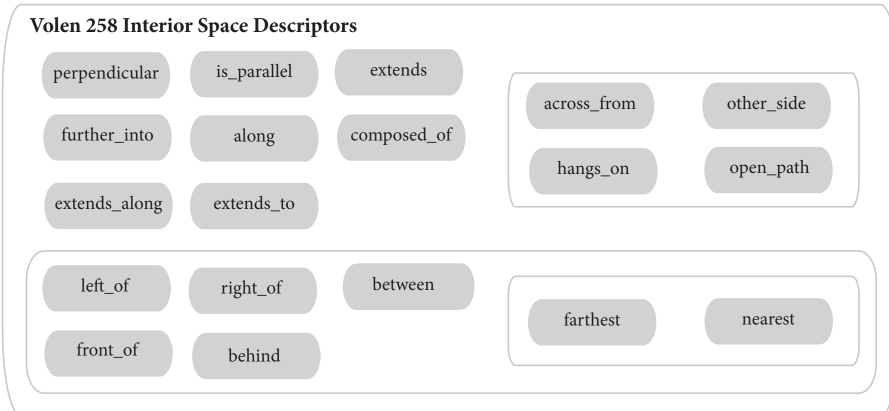

FIGURE 9.2 Key Spatial Descriptors

The aim of this chapter is to: (a) pragmatically illustrate the manner in which spatio-linguistic conceptions occur within a range of spatial assistance systems involving human assistance, assurance, and empowerment; (b) provide an abstraction and formalization mechanism that may be used as an interfacing mechanism between the spatio-linguistic conceptualization of humans, and their precise interpretations within a practical system; (c) with respect to our perspective on spatial information handling for a stated purpose, illustrate the nature, role, and significance of computing spatial relations within practical spatial assistance systems.

The chapter is organized as follows: Section 9.2 lays out our interpretation of a spatial assistance system. Section 9.3 focuses on select application areas from the range introduced in our general discussion of spatial assistance in Section 9.2. Section 9.4 develops the formal framework that is required to operationalize the abstractly identifiable notion of the structural form (also developed in this section). The crux of this section lies in the elaborations on multi-perspective semantics and qualitative abstraction mechanisms in the context of structural forms. Section 9.5 demonstrates the spatial computing and underlying reasoning patterns that may be realized with the formal model presented in the chapter. Finally, Section 9.6 includes a brief summary of the chapter, together with a discussion of broad perspectives and aspects not covered in the chapter in detail.

9.2 What is a Spatial Assistance System?

A Spatial Assistance System (SAS) is a computational manifestation of the spatial decision-making and other forms of analytical and creative abilities situated along a multi-dimensional and mutually interacting spectrum: on the one hand are those abilities that typically require extensive domain-specific training, knowledge, and expertise (e.g. an architect, a film cinematographer), on the other are those abilities that merely require the natural intelligence that humans are equipped with by virtue of their everyday existence (e.g. wayfinding tasks). Regardless of the precise domain of application, the crucial developmental aim of a spatial assistance system is to transfer the cognitive stress involved in a human analytical activity onto a system, by externalizing and operationalizing the decision-making processes involved therein. In essence, spatial assistance systems are basically instruments of human *assistance*, *assurance*, and *empowerment* and they serve one or more of these functions depending on the precise area of their application.

This chapter is focused on computational systems and frameworks for spatial awareness capabilities comparable to those of humans. Some instances of SASs include decision-support tools that require specialized *spatial reasoning capabilities*, for instance, in the manner as defined for the domain of intelligent spatial design assistance (Bhatt and Freska, 2010). As a broad range of examples, consider the following application areas where the aforestated notion of spatial assistance may be clearly identified:

- *Spatial design assistance.* The domain of spatial design assistance encompasses all those areas where humans engage in a creative spatial design or planning activity of some form. As a key example, consider the case of architectural design assistance systems where a work-in-progress design of a building within the context of a Computer-Aided Architecture Design (CAAD) tool has to be reasoned about. Other scenarios include urban planning, tasks involving spatial conceptualization and visual thinking about spatial structures, for example pre-production assistance in media design.
- *Real-time emergency assistance.* These are systems that provide intelligent assistance for emergency scenarios such as fire-fighting, rescue, and evacuation, and paramedic-support in emergency scenarios. This domain is characterized by the need to provide high-level strategic-planning assistance, for example in order to meet the immediate challenges of navigating and overcoming physical obstacles in an unfamiliar and hostile (e.g. low-visibility) environment.
- *Ambient intelligence (AmI), pervasive, and ubiquitous Computing (UbiComp).* These include a broad range of technologically driven systems involving the modelling of, for instance, action, change, interaction, situational context, and the semantics of *space* and *place* within practical deployments in the real world. Some instances include mobile and location-based services, systems of activity modelling, and behaviour interpretation and analysis.
- *Indoor navigation assistance.* These are systems that provide specialized wayfinding and navigation support mechanisms for built-up environments such as

malls, exhibition centres, museums, airports, train-stations and other categories of built-up public spaces.
- *Ambient Assisted Living (AAL).* This is a class of practical systems aimed at serving an empowering and assuring function within private spaces such as residences and offices. Typically, these systems involve interactions between humans, robots, and software systems.

Whereas there exist distinct categories of spatial assistance systems differing in the domain of application and the precise manner of intelligent assistance they provide, there are several fundamental similarities underlying the systemic and information-theoretic aspects of these systems. Primarily, the similarities pertain to the ontological, representational, and computational foundations that underlie their practical design and implementation. Specifically, central to these categories of assistance systems is a common foundational basis consisting of representational modalities and computational capabilities:

- from a *representational viewpoint*, modalities for semantic modelling, multi-perpective representations, and qualitative spatial abstractions acquire a central significance;
- from a *computational viewpoint* and closely connected to the representational modalities, computational techniques for conceptual and spatial reasoning define the essential character and nature of the (spatial) analytical and assistive capability that is implemented.

In essence, one may identify several fundamental capabilities with respect to the spatial conceptualization, modelling, and computing capabilities within the context of the range of systems identified herein. For instance, consider the case of spatial design assistance and emergency and navigation assistance systems. The information and computational requirements for *spatial reasoning* in the context of an indoor *Structured Spatial Environment* (SSE) bears close relationships and similarities[4]—several capabilities, for example by way of wayfinding complexity analysis, (real-time) wayfinding assistance, data analysis, and artefactual simulation, virtual reality, natural human–computer interaction that may be operationalized within these systems can be easily identified. Furthermore, the conceptualization and representation of quantitive descriptions of real (i.e. already existing) or hypothetical (i.e. being designed) indoor spatial environments is also based on shared foundations. For the case of indoor or built-up environments, and for spatial assistance scenarios such as those aforementioned, it may be presumed that geometric model(s) of the environment under consideration are available, for example by way of accurate building and floor

[4] These models of SSEs need to be grounded to industrial data representation standards designed for community-wide tool compliance and interoperability. Relevant remarks concerning this aspect may by found in Section 9.6.

plans (CAAD, design assistance), graph-based models (wayfinding assistance), and finite-element models based on computation fluid dynamics (for structural analysis, cost estimation, phenomenal studies to simulate fire spread). These models may pertain to real spatial environments that have been built (e.g. a museum), or they may pertain to an arbitrary environment that is undergoing initial conceptualization, prototyping, and design.

Spatial reasoning (for spatial awareness), however it may be defined from a cognitive, ontological, and computational viewpoint, does not differentiate between real and hypothetical environments. That is, different types of analytical capabilities that may be deemed to be within the purview of a particular interpretation of spatial awareness have to be based on high-level quantitative and qualitative perspectives that are grounded to a geometric model of an environment that may exist either in reality, or merely as a hypothetical construction within a system.

9.3 The spatio-linguistic markers within Spatial Assistance Systems: select case studies

The aim of this section is to present high-level, yet concrete examples of the nature of spatial assistance that is applicable within a select category of spatial assistance systems. The approach here is to illustrate the spatio-linguistic conceptions that accrue within the scope of our selected case studies. These case studies are then further elaborated on in the rest of the chapter.

9.3.1 Architecture design assistance

Spatial design in general, and architectural design in particular, as a problem-solving activity typically consist of the conception—modelling—evaluation—re-modelling cycle. Essentially, a designer in this case is considered to be an *agent of change*, who may be intuitively regarded as traversing a complex configuration space of possibilities, and selecting one course of action (guided by domain knowledge, expertise, cognitive capabilities, specialized requirements, aesthetic preferences, and so forth) that produces a desired product or design. Since architectural design tasks are concerned with a spatial environment, formal representation and reasoning along conceptual and spatial dimensions is essential to ensure that the designed model satisfies key requirements that enable and facilitate its intended function.

9.3.1.1 A design task As a use case, consider the task of initial conception and design of a museum. A museum is an instance of a *structured spatial environment* that not only has a desired form and function, but is also constructed keeping in mind pre-determined aesthetic, cultural, psychological, and other subjective parameters. For example, consider the high-level spatio-linguistic conceptualizations of sobriety, austerity, and comfort with respect to low-level functional features such as connectivity, spatial distribution, and organization:

the larger trees are located in an area that is more elevated than the whole northern rim of the plot . . . covering creates a gentle artificial elevation that perspectively accentuates and enhances the whole architectural composition. The distribution of the construction volumes fundamentally followed a desire for horizontality, allowing one to read the continuity of the green space beyond the construction and in all directions.
This sober, rational and markedly horizontal structure is distinctive for its laminar exterior, its modular repetition, austere design and the hard quality of the materials which shape it, concrete and glass.
The main matrix-like structure had long units and was spatially austere, linking public and private spaces while at the same time delineating the two inner patios of the Museum. (Museum Gulbenkian, Lisbon. Architecture and Landscape
(Tostoes et al., 2006: 22, 23, 26, 27))

9.3.1.2 Statutory building codes In addition to high-level conceptions of the design, designers are confronted with mandatory regulations that enforce several structural constraints at different levels of complexity:

Steps of a staircase may not be connected directly to a door that opens in the direction of the steps. There has to be a landing between the staircase steps and the door. The length of this landing has to have at least the size of the door width. (Staircase / Treppen (§35(10), pg. 24. Bremen (Germany) Building code (BremLBO, 2003))

9.3.1.3 Design guides In addition to the statutory codes, one may also identify design requirements emanating from expert recommendations. For instance, the following (statutory or optional) requirements may be identified from the US Courts design guide (US GSA, 2007), and Alexander's *pattern language* (Alexander et al., 1977):

Courtroom areas used by the public must be accessible to people with disabilities. Private work areas, including the judge's bench and the courtroom deputy, law clerk, bailiff, and court reporter stations, must be adaptable to accessibility. While all judges' benches and courtroom personnel stations do not need to be immediately accessible, disabled judges and court personnel must be accommodated.

All architectural elements must be proportional and arranged hierarchically to signify orderliness. The materials employed must be consistently applied, be natural and regional in origin, be durable, and invoke a sense of permanence (US Courts Design Guide 2007 (US GSA, 2007))

Place the main part of the kitchen counter on the south and southeast side of the kitchen, with big windows around it, so that sun can flood in and fill the kitchen with yellow light both morning and afternoon. (Alexander et al., 1977)

Environmental feature descriptions such as the ones here mentioned refer to abstract, high-level spatial design patterns that correspond to specific structures at a quantitative level. For instance, it is noticeable from the descriptions of the museum that early design and conceptualization involved high-level feature descriptions of the *structural form* of the environment. Spatial features such as *continuity*, *spaciousness*, *symmetry*,[5] *modular repetition*, *elevation*, *relative positioning of entities*, *visibility relationships*, *(barrier-free)*, and *accessibility* may be easily identified. Section 11.4.2 will further illustrate the manner in which such high-level features may be modelled and reasoned about using the abstraction mechanisms that we propose in this chapter.

Contemporary professional design tools, and the precise quantitative modelling paradigm that they are based on, are incapable of exploiting the correspondence between high-level descriptions of spatial concepts and features. Such tools simply lack the ability to exploit the expertise that a designer is equipped with, but is unable to communicate to the design tool explicitly in a manner that is consistent with its inherent human-centred conceptualization, that is, *semantically* and *qualitatively* (Bhatt and Freska, 2010). This chapter, in its remaining parts, illustrates the manner of formalization and computation that is needed to develop the assistance capability to solve design problems of this nature.

9.3.2 Creative assistance in media design

We interpret *creative assistance in media production* as the capability of computational tools to augment the creative capabilities of experts and artists such as cinematographers, directors, script and screenplay writers, and storyboarding artists at several stages within the media design and creation process. Consider the domain of film and comic book pre-production. Here, one may identify several forms of assistance at the production phase, for example virtual cinematography, storyboarding, and scene visualization from scripts and automatic camera control in the animation domain. Some examples from domains of our active interest follow (Bhatt and Flanagan, 2010).

9.3.2.1 Scenario description in media pre-production As an example, consider a typical creative design process between a script or screenplay writer, and storyboarding artist and cinematographer or director. The illustration in Figure 9.3 is a snapshot from the freely available media pre-production software CelTx, which is designed from the viewpoint of a wide range of design domains. It facilitates creation and organization of media projects like screenplays, films, videos, stageplays, audio plays, documentaries, machinima, comics, games, and podcasts.[6]

The illustration of Figure 9.3 corresponds to parts of a script and screenplay for a comic strip, as re-produced below:

[5] For example, one way to ensure a sense of '*justice*' and '*fairness*' within a court-room is to ensure a symmetric balance and equal access to (elements from) both sides of the courthouse.

[6] Celtx. A Media Pre-Production Environment. www.celtx.com/

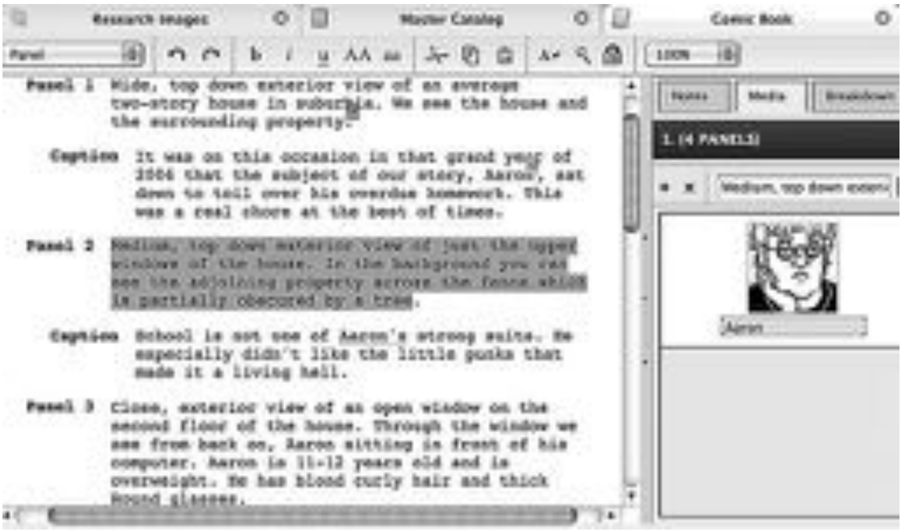

FIGURE 9.3 CelTx: A (free) Media Pre-Production Software.

Scene: Wide, top down exterior view of an average two-story house in suburbia. We see the house and the surrounding property.
Scene: Medium, top down exterior view of just the upper windows of the house. In the background you can see the adjoining property across the fence which is partially obscured by a tree.
Scene: Close, exterior view of an open window on the second floor of the house. Through the window we see from back on, Aaron is [. . .]
Scene: A wide panel that uses the lower half of the page. A close-up view of Aaron. He is staring directly ahead [. . .]

(The Mechanical Shakespeare; Text source: Comic Strip, Celtx Pre-production Software.)

As a basic level of assistance, the capability to automatically produce scenario visualizations based on the semantics of the spatial content in the discourse material is useful. For instance, it is typical for cinematographers and set designers to use the skills of storyboarding artists to start conceptualizing the precise manner and technical method of actually filming or animating a sequence. Consider the example storyboards included in Figures 9.4(a)[7] and 11.4(b).[8] Here, one may imagine the automatic production of such storyboards on the basis of the *structural form* that is semantically interpreted and derived from the scenario description of the scene.

[7] Art credit: Wikimedia, Wikiversity. www.wikiversity.org/
[8] Art credit: Peter Rubin; Goldman et al., 2006.

Thumbnail Storuboard - page 2

(8) Old Person: “What computer do you have at home?”

(9) Young Person (eagerly): “A Macintosh!”

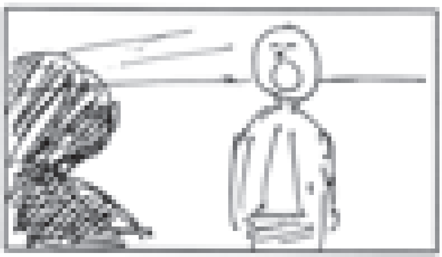

(10) Old Person: “But what computer does your father use at work?”

(11) The young person thinks for a moment. Young Person: “Humm!”

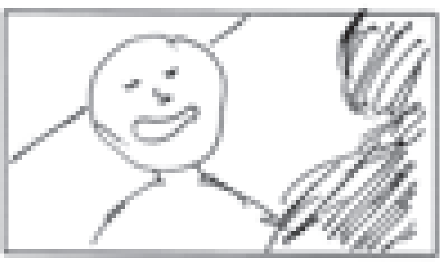

(12) Young Person (amazed and excited): “Seduced by the Dark Side!”

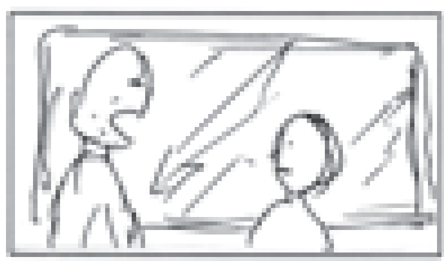

(13) The old person smiles. Old Person: “Ahh!”

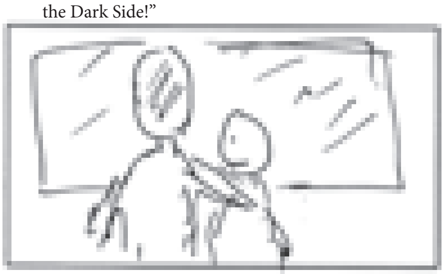

(14) The old person and the young person start to walk toward home.

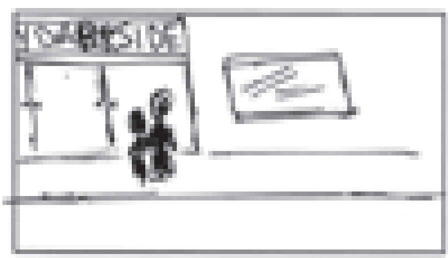

(15) The old person and the young person walk toward home together.

(a) A Storyboard

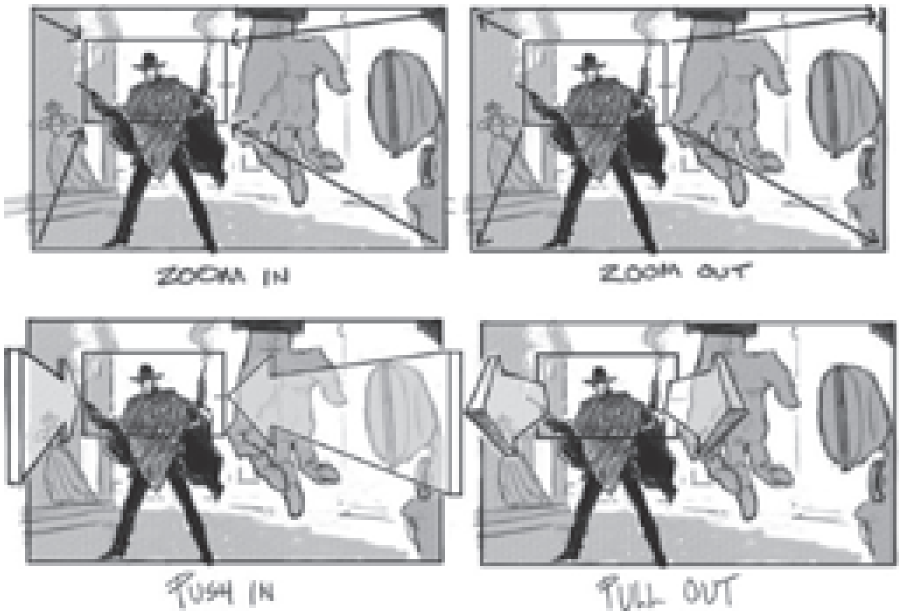

(b) Action-Annotated Storyboards

FIGURE 9.4 Automatic Storyboarding

The generation of these storyboards is connected to the much broader goal task of scenario and narrative completion, as explained below.

9.3.2.2 Scenario and narrative completion In general, the field of automatic cinematography aims to derive a sequence of camera shots (i.e. the camera's orientation to the actors, camera's focus, angle of view, etc.) from descriptions provided in a script (Drucker and Zeltzer, 1995; Christianson et al., 1996; He et al., 1996; Lu and Zhang, 2002; Bhatt and Flanagan 2010). Most automatic cinematography involves using a knowledge base of filming heuristics to control the perspective or placement of a camera based on contextual cues of the scene. In this context, a film can be viewed as a hierarchy (He et al., 1996); the top of the film hierarchy is the script, which consists of a sequence of time-ordered narrative descriptions (Bhatt and Flanagan, 2010), referred to as scenes. Each scene, in turn, provides contextual information, in the form of actions and events that can be used to derive a specific camera shot. The objective of each camera shot is to capture the sequence of events in a manner that is cinematically 'pleasing', that is, it achieves a pre-determined aesthetic, dramatic, or emotional effect.[9]

[9] For elaborate studies on the theory, art, and semiotics of narration, and the 'language games' that underlie this domain, the collected works of Edward Branigan (1984, 2006) should be investigated.

(a) Establishing Shot

(b) External Shot

(c) Close-up / Reaction Shot

FIGURE 9.5 '*Film Idioms*' are heuristics for determining the relative positioning of actors and the camera within a scene

As an example, consider the simple but common scene in Figure 9.5[10] depicting a group of two actors. Within this scenario, the context of each scene is based on the current state of each actor with regards to their participation in the conversation, that is *talking*, *listening*, or *reacting*. Below is a sample script that involves two actors,

[10] Stills credit: '*Mesrine: L'instinct de mort*' (2008), France. Director: Jean-Francois Richet. www.mesrinemovie.com/

Kendra and Annika, engaged in a conversation. In the example, contextual cues are provided as key words that indicate the current state of each actor: 'Kendra starts to *talk*' and 'Annika *reacts* with astonishment' and so forth, whereas from the screenplay writer's perspective, the manner in which the scene has been conceptualized is based on heuristics guiding the placement of the actors and entities in relation to the location of the camera:

Act: Kendra and Annika
[Establishing-shot] – Kendra and Annika
Kendra starts talking to Annika – ["*dialogue*"]
[Cut: mid-shot] – Annika reacts anxiously to Kendra
Kendra continues talking to Annika
[Cut: Close-up] Annika responds to Kendra – ["*astonishment*"]
End.

(A Sample Narrative)

As the scenes progress and the conversation develops, the states of the actors change. From this information, it is the job of the (automatic) cinematographer to decide on an appropriate sequence of camera shots to properly depict the conversation. The result of this process is similar to the storyboard found in Figure 9.5, or the virtualization depicted in Figure 9.6. They show the perspective of the camera throughout the key moments of the scene. Because this scenario is so common in film, cinematic patterns have emerged that define heuristics to capture this particular type of situation, referred to by cinematographers as a *film idiom* (Arijon, 1976). These idioms have been defined for many typical cinematic situations, such as groups of actors in a conversation, or an action sequence. For instance, the 'spatial structure' associated with a film idiom may be formalized using qualitative spatial abstraction in a manner depicted in Figure 9.6. In general, a film idiom can be seen as a set of declarative rules that specify a mapping between the use of camera shots to a situational context. We formally build up on these aspects in the sections to follow, and in Section 9.5, illustrate the nature of spatio-temporal computation that is necessary to perform reasoning (e.g. scenario consistency, scene interpolation, or spatio-temporal abduction) in this particular domain of interest.

9.3.3 Real-time emergency assistance

The domain of emergency rescue work is characterized by urgent, safety-critical decision-making based on very limited and vague information that is dynamic, volatile, and possibly erroneous. These challenges exist at a range of abstraction levels, from the high-level strategic planning of firefighting captains to the immediate challenges of navigating and overcoming physical obstacles in extremely hostile, noisy, and low-visibility environments faced by firefighters. Navigation requires a

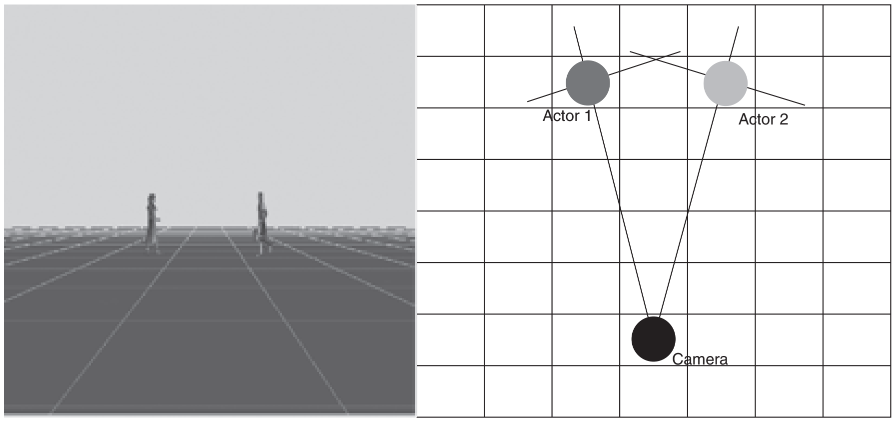

(a) Establishing Shot

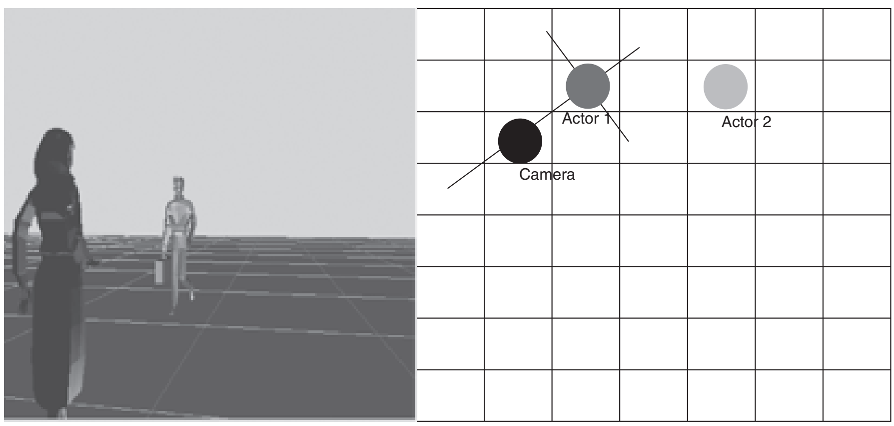

(b) External Shot

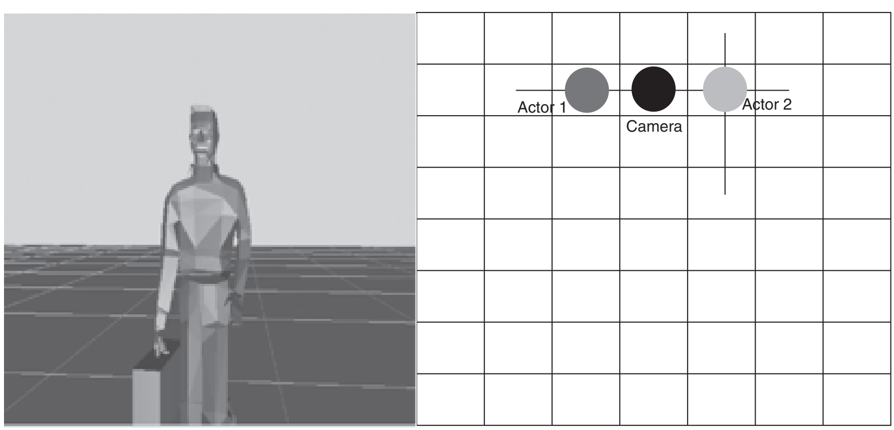

(c) Reaction Shot

Figure 9.6 Minimal structural form of film idioms modelled using ternary orientation relationships. Virtual cinematography; *2 Avatars and 1 Virtual Camera*; animation shows perspective of the black circle, which is a virtual camera.

combination of spatial modalities; in particular, the standard definition of route graphs (Werner et al., 2000) that is based on the connectivity of regions is effective at the geographic level, but is too general for providing the spatial references and directions needed by firefighters as they navigate through a burning building with zero visibility. Instead, a more specialized definition of route graphs is required that is modelled on the spatial language that firefighters use when communicating and navigating through hazardous indoor environments. For example, consider two firefighters navigating through one of the buildings at a power station site in search of victims (Figure 9.7); the following communication takes place:

FFB: [. . .] we're standing by the door where we came in . . so now we can put the right hand on the wall here

⋮

FFA: do you feel a wall on the right side?

FFB: yeah, I've had contact with the wall the whole time there was a shoe shelf on the right side so

⋮

FFB: <[. . .] we've gone into the reference room again, so now we're going into the second door on the left side>

FC: <ah, you went into door two I heard, over>

⋮

FFA: [. . .] ok . . then we have one room left in the first room, there was a door, we haven't been in there

⋮

FFA: <[. . .] we went straight on the first door on the left side and searched it>

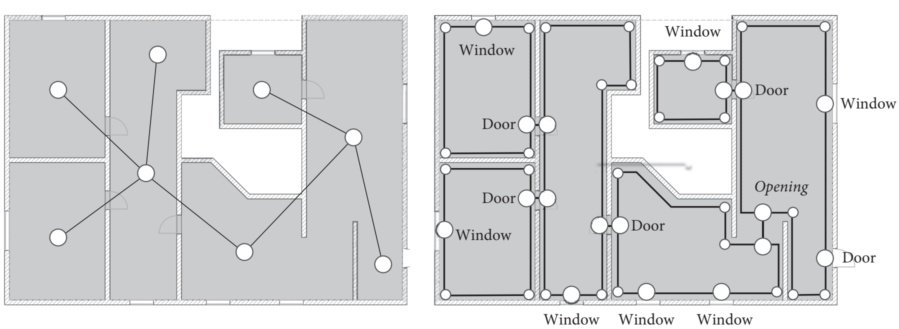

(a) Region-connectivity route graph of a building in the power-station complex.

(b) Specialised route graph based on the positioning of feature along walls of rooms in a building in the power-station complex.

Figure 9.7 Specialised route graphs provide a more effective modality that supports firefighters' current approaches to navigation compared to the standard definition of route graphs (floor plan adapted from the ground floor of the Stockholm Fire Department training facility in the Agesta training centre, as cited in [39, 40]).

FC: <eh, can you repeat, over>
FFA: <[. . .] we took the left hand on the wall and [. . .] went into the first room>
⋮
FFA: [. . .] I've reached a doorway on the left side and we are going in there and we're holding the right hand on the wall [. . .] and see if we find the fire-extinguisher, over

(communication recorded during firefighter training exercise by Lindgen, 2004; Lindgren et al., 2007)

As is clear from the above communication, the firefighters' sense of orientation depends heavily on reference features such as doors, walls, corners, and large pieces of furniture. As illustrated in Figure 7(a), the standard route graph does not provide the type of information that a firefighter needs when navigating through a building. A more effective, domain-specific route graph is defined by the arrangement of salient features such as doors and windows along room walls, as illustrated in Figure 7(b).

The chapter by Gallay et al. (this volume) focuses on devices that provide navigation assistance for blind pedestrians. Within this context, the task we present for first response emergency rescuers can be characterized as requiring global navigation assistance via either visual or non-visual interfaces in a low-to-zero visibility environment; as a contrast to the chapter by Gallay et al., we focus on assistance through a visual interface.[11] For example, consider the environment from a firefighter's perspective as they enter the building, as illustrated in Figure 9.8. A navigation assistance system is mounted on each of the firefighter's helmets with a small transparent display on the mask. The mask display lists the features of the room that are useful for orientation, ordered using the specialized route graph in Figure 7(b). The standard route graph can be used when the firefighters can move freely through the room, although when visibility is significantly reduced, the specialized route graph is more effective for navigation, as the firefighters need to rely on walls for orientation.

The role of (assistive) spatial reasoning is to augment the firefighters' sense of the surrounding space and enhance their comprehension of the environment, for example, by informing them about local reference objects and their relative qualitative spatial relationships. Effective information technology tools can assist in strategizing about search routes within the building, taking into account temperature sensor readings, fire spread models based on the known characteristics of the incident floor plans of the building, and so on.

[11] In contrast to pedestrian navigation, the relatively fine-grained localization required for indoor navigation can be achieved using a combination of dead reckoning strategies (Miller, 2006).

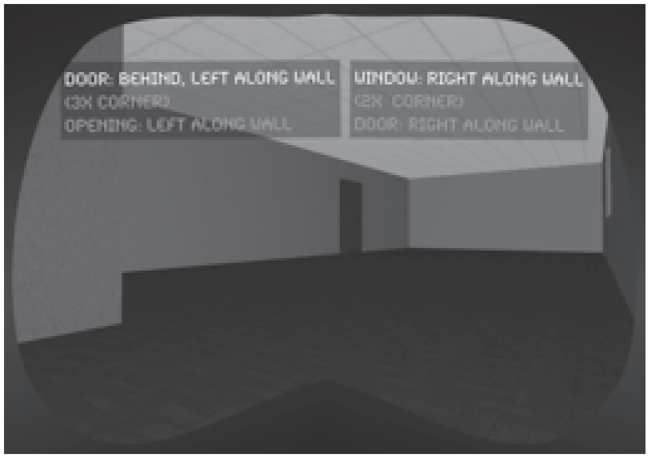

(a) Firefighter perspective with no smoke; the standard route graph is applicable for providing navigation assistance.

(b) Firefighter perspective in burning building with smoke-filled interiors; the specialsied route graph is required for navigation assistance.

FIGURE 9.8 Navigation assistance as firefighters enter the burning building.

9.4 Structural form: multi-perspective representational semantics and modal abstraction

Having exemplified the spatio-linguistic markers that are identifiable as a point of human and spatial assistance system interaction in Section 9.3, the aim in this section is to elaborate on the internalization mechanisms of the respective linguistic markers within the system. The section situates itself in the context of the range of spatial assistance systems introduced in Section 9.3.

9.4.1. Formalizing spatial structure

Reconsider the range of spatial assistance systems introduced so far in Sections 9.3.1–9.3.3 by way of the architecture design assistance, creative media design assistance, and the real-time emergency assistance domains. In each of these domains, a human stakeholder is involved in the spatio-linguistic conceptualization and its communication with a system using some human–computer interaction modality. In essence, there exists some human-centred conceptualization of the spatial structure of a real or hypothetical environment, scene, or more abstractly, a *structural form*.

The following abstract notion of the *structural form* of an environment is identifiable:

> The *structural form* of an environment is an abstraction generally corresponding to the layout, shape, relative arrangement, composition at the common-sense level, of spatial entities, artefacts, and anything else—abstract or real—that may be geometrically modelled, interpreted, or derived within a design system. The only conceivable premise underlying this notion is that it should be possible to communicate the conception of the structural form using one or more spatio-linguistic modalities—e.g., spatial prepositions, path and pattern descriptions, region and point-based abstractions—that may be wholly or partially grounded to an underlying physical structure either in metric space, or in an abstract qualitative space.

For instance, the structural form may be minimally interpreted as a constraint network that determines the relative qualitative spatial relationships between the real and artefactual entities (Section 9.4.3; Figure 9.10) contained within a design. A *scene description matrix*, that is., a two-dimensional table characterizing the spatial relationship of each entity with every other entity in the model, could be the minimal basis of qualitatively abstracting an indoor spatial model. Indeed, from a formal modelling viewpoint, the qualitative model would be based on a semantics that is cognitively and linguistically grounded, and conforms to the formal relational semantics of a spatial logic (Aiello et al., 2007).

The above-stated notion of structural form based on a complete scene model is indeed minimal, and several possibilities exist for further refinement. For instance, as opposed to a complete scene description that characterizes the spatial relationship of an entity with every other entity, one may resort to a hierarchical model that exploits the natural order of organization in the physical environment. Similarly, hierarchies may be generated not only based on physical structure, but also on the basis of semantic organization. It is useful to be able to characterize *spatial patterns* such as 'circularity', 'extends_along', 'further_into', 'open_path', 'composed_of', etc. that were identified in the context of the Volven 258 scenario illustrated in Figure 11.1 in Section 11.1.

9.4.2 Multi-perspective semantics and representational modularity

Given the interpretation of structural form, an abstraction such as a 'Room' or 'ArchitecturalEntity' may be identified semantically by its placement within an ontological hierarchy and its relationships with other conceptual categories. This is what a designer must deal with during the initial design conceptualization phase. However, when these notions are transferred to a tool (e.g. a CAAD or a media pre-production tool), the same concepts acquire a new perspective. In the case of the CAAD tool, the designer must deal with points, line-segments, polygons, and other geometric primitives available within the feature hierarchy of the design tool, which, albeit necessary, are in conflict with the mental image and qualitative conceptualization of the designer. Likewise, a 'Floor' at the designer's conceptual level is abstracted as a 'Region' at the qualitative level of a reasoner and as a 'ClosedPolygon' within the geometric perspective of the design tool, thereby preserving the geometry at the quantitative level of a CAAD-based feature model (Figure 9.9). Multi-perspective representational semantics enables a knowledge-based system to make inferences about the conceptual design and its geometric interpretation within a CAAD model in a unified manner. On the representational front, the concept of multi-perspective semantics may be readily translated to representational modularity, as identified within the ontology / knowledge engineering, and conceptual modelling disciplines (Stuckenschmidt et al., 2009).

9.4.3 Multi-modal abstractions

This abstract view of spatial structure can be grounded to reality via the medium of modalities, namely, semantic, qualitative, and other forms of (graph-theoretic) abstractions that serve as an interface between the spatio-linguistic conceptualization of structural forms, and their concrete interpretation within a spatial assistance system. The following abstractions may be identified.

9.4.3.1 Spatial artefacts Semantic descriptions of designs and their requirements acquire real significance when the spatial and functional constraints are among strictly

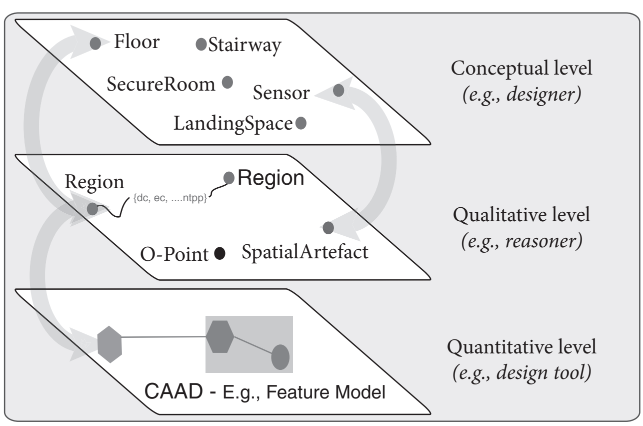

Figure 9.9 Multi-Perspective Representational Semantics

spatial entities as well as abstract *spatial artefacts*. For instance, it is possible to model the spatial layout of an environment at a fine-grained level; but it is not possible to model spatial artefacts such as the range space of a sensory device (e.g. camera, motion sensor, viewpoint of an agent) in the same way. Spatial artefacts do not necessarily have a material existence; nevertheless they need to be treated as such. In general, architectural working designs only contain physical entities. Therefore, it becomes impossible for a designer to model constraints involving spatial artefacts at the design level. For instance, consider the following constraint: '*the motion-sensor should be placed such that the door connecting room A and room B is always within the sensor's range space*'. The following spatial artefacts may be identified (Figure 9.10).[12]

A1 the *operational space* denotes the region of space that an object requires to perform its intrinsic function that characterizes its utility or purpose;
A2 the *functional space* of an object denotes the region of space within which an agent must be located to manipulate or physically interact with a given object;
A3 the *range space* denotes the region of space that lies within the scope of a sensory device such as a motion or temperature sensor, or any other entity capable of visual perception. Range space may be further classified into other categories, such as *observational space* (e.g. to model the concept of the *isovist*[13]).

9.4.3.2 QvGraphs We propose Qualitatively Annotated Visibility Graphs (QvGraphs) as an extension to the concept of a *Visibility Graph* (Lozano-Pérezard Wesley, 1979; de Berg et al., 2000). In computational geometry, a visibility graph of a polygonal scene shows the intervisibility relations between a set of points (indicating locations, obstacles, and so on) in a scene, as geometrically constituted within the Euclidean plane. Specifically, visibility graph vertices correspond to point locations and edges represent a visible connection between them. QvGraphs extend visibility graphs by deriving and annotating the *visibility link* with (potentially disjunctive) knowledge about spatial relationships pertaining to one or more spatial domains such as topology, orientation, and distance. Figure 9.11(a) illustrates an example of a visibility graph of a museum lobby. The direction of the edges indicates the direction of the binary qualitative relations; for example, the 'ReceptionDesk' is 'right_of the LobbyEntrance', indicated by the direction of the edge in the QvGraph, although the 'visible' relation in this example is symmetric.

9.4.3.3 Spatial sequence graphs In natural language, it is common to refer to a sequence of objects, where the objects are ordered along some path through the environment. Consider the following expressions:

[12] Formal definitions of spatial artefacts may be found in Bhatt et al. (2009).
[13] An isovist is the set of all points visible from a given vantage point in space and with respect to an environment (Benedikt, 1979).

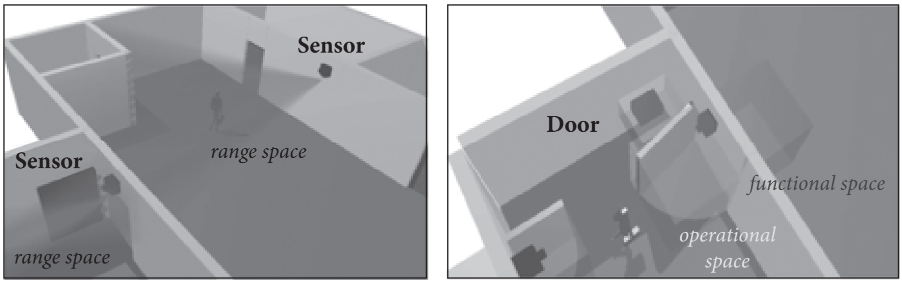

(a) Implicit artefacts within a design.

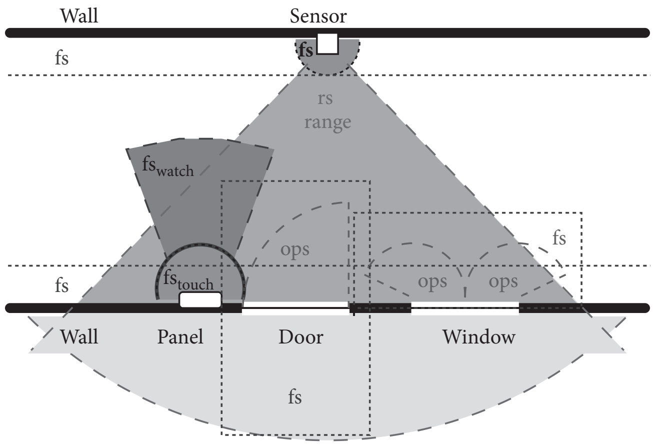

(b) Floor plan perspective of the implicit artefacts within a design.

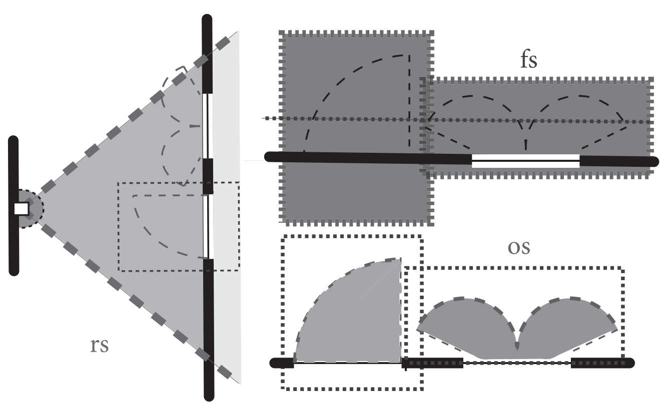

(c) Range space (rs), Functional space (fs), Operational space (os)

FIGURE 9.10 Spatial artefacts are entities, which unlike regular spatial objects, do not have a physical manifestation in reality (or within a design), but need to be treated as such for all practical / reasoning purposes.

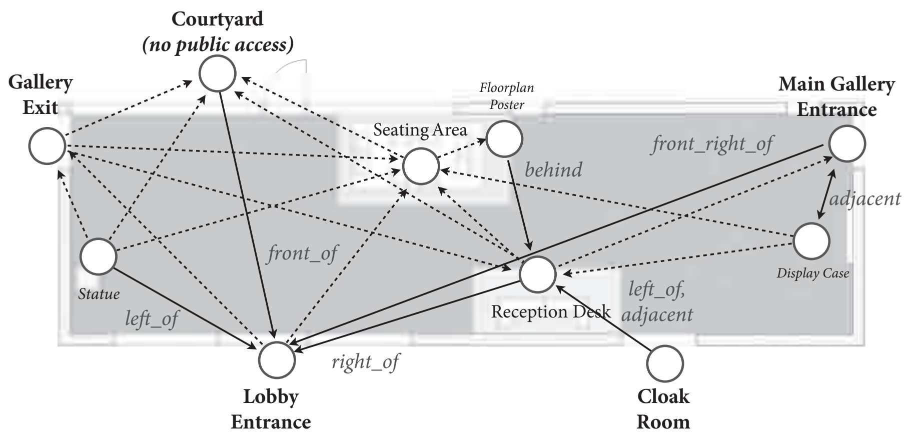

(a) Partially annotated QvGraph of the lobby. The user has specified that orientation and topological relations are relevant for this QvGraph (the qualitative annotations on the dash edges have been omitted for clarity).

(b) The Real Environment: Lobby Area

FIGURE 9.11 Lobby Area at Museum Gulbenkian – QvGraph Analyses

- *Numerous paintings are mounted along the wall.*
- *The Far East art section is down the room after the Oriental-Islamic and Armenian rooms.*
- *Rivets have been placed evenly along the edge of the column.*
- *Further into the room is a group of partitions.*

In each of these examples a virtual *path* has been implicitly defined, and the objects have been partially or totally ordered along this path. The paths typically follow the shape of some reference object such as a wall, beam, table surface edge, and so on. Moreover, the path is *directed* giving meaning to the terms *before* and *after*; one example is by specifying the start of the path to be the object that is nearest to the person referring to the sequence of objects. Note also that paths may be a simple cycle consisting of a loop involving all objects, for example, art pieces positioned along the complete perimeter of a gallery room.

This notion is formalized as a *spatial sequence model* where vertices represent objects and directed edges represent the object ordering. Edges are optionally annotated with any useful additional qualitative spatial relations between the ordered objects. Figure 9.12(a) illustrates an example of two spatial sequence models in one of the Gulbenkian Museum gallery rooms.

9.4.3.4 Route graphs A route graph, as defined in Werner et al. (2000), corresponds to a cognitively and linguistically motivated spatial representation of an environment that focuses on qualitatively capturing different routes an agent can use for navigation. The standard definition of route graphs is based on the connectivity of spaces (rooms for example), such that an agent can move freely from one space to another without necessarily passing through an intermediate space. For example, Figure 9.13 illustrates the route graph (from the perspective of art gallery visitors) of the entire Gulbenkian floor plan, and Figure 9.14 illustrates the route graph of the Gulbenkian lobby.

There is a strong connection between properties of structural layout (such as regularity), floor plan complexity, and properties of navigability that a building design affords (i.e. separate from other means of navigation such as signage) (Peponis et al., 1990; O'Neill 1991a, 1991b; Baskaya et al., 2004; Werner and Schindler, 2004. Route graph analysis can greatly assist in tasks such as building design and navigation when the route graph is derived from the appropriate structural properties according to the particular application domain. That is, domain-specific specializations of connectivity are required to model movement in different applications. Moreover, these heterogeneous route graphs must be integrated in a manner that corresponds to a person's local and more global topological comprehension of an environment (Haq and Zimring, 2003).

9.4.3.5 Flow vectors The topological information represented in route graphs is not rich enough (or at least does not make the necessary region distinctions) to specify certain qualitatively significant movement patterns of people and objects, such as modelling airflow in a relatively confined room connected to the building ventilation system (Kowadlo and Russell, 2006). Such movement patterns cannot be sufficiently expressed using route graphs without first introducing new approaches for partitioning a room into regions that are only relevant for adequately modelling

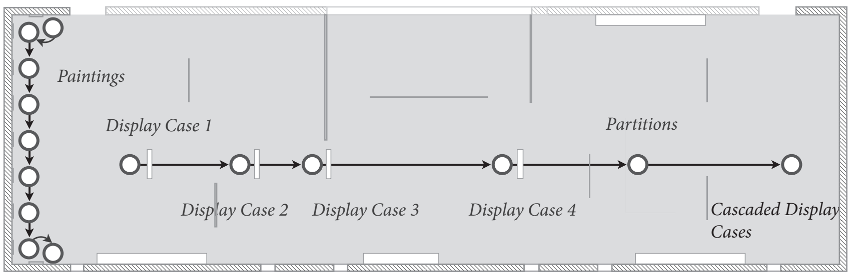

(a) Two spatial sequence models

(b) The Real Environment

FIGURE 9.12 Oriental-Islamic and Armenian Gallery Rooms at Museum Gulbenkian – Spatial Sequence Graph Analyses

some particular movement phenomenon. Rather than introducing numerous specialized region distinctions, these distinctions can be implicitly embedded in the definition of a new type of model called the *flow vector graph*. Flow vector graphs are derived by directly focusing on the physical movement patterns of agents and objects rather than on the a priori definition of connectedness of the spaces that the agents and objects are moving through. Flow vector graphs are closely tied to the underlying geometry of spaces and the semantics of objects. That is, rules for deriving flow vector graphs can specify different movement patterns depending on

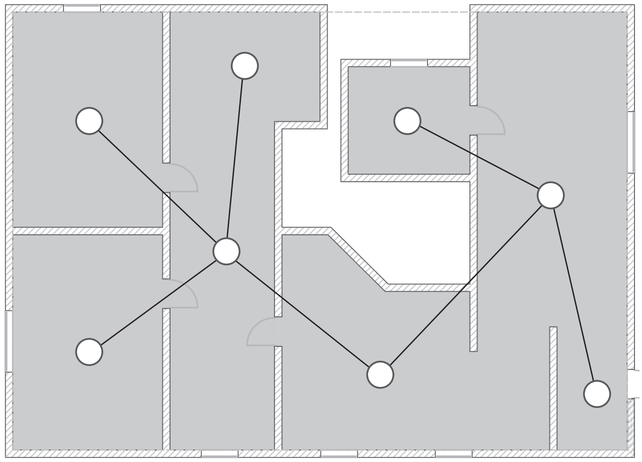

FIGURE 9.13 Route graph of the Gulbenkian floor plan from the perspective of visitors.

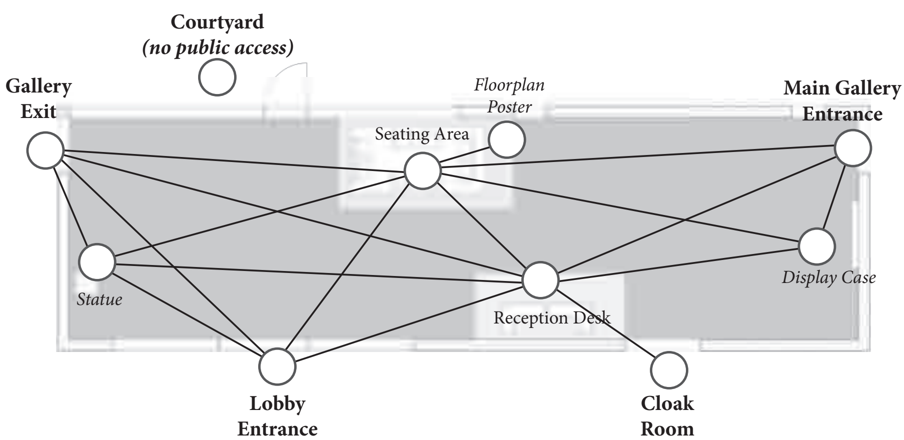

FIGURE 9.14 Route graph of the Gulbenkian lobby from the perspective of visitors.

whether an object is a *statue* or a *chair*. As with route graphs, flow vector graphs typically either specify movement between different spaces within a building or specify local movement within a space (such as a room).[14]

9.4.3.6 Hierarchical models The data access framework provides access to a *hierarchical* and *multi-domain* model of space. From the viewpoint of hierarchization, the aim of this work is to develop an organization of qualitative spatial information that

[14] Flow vector graphs are further elaborated on with an example in Section 9.5.3; also see Figure 9.18.

splits the related entities into independent subsets and allows for solving spatial reasoning tasks at an adequate level of granularity. The resulting hierarchical representation should support the same reasoning and design tasks that would be possible with a flat qualitative representation but do so in a more efficient and intuitive way.

Hierarchization of buildings is defined based on spatial containment and spatial aggregation, such that a higher-level feature contains or consists of the lower-level features in the hierarchy. These hierarchical models can be based on either structural or semantic relationships. Structural containment refers to the strictly physical aggregation and containment of the components of a building; for example, a building is composed of floors (or storeys), each floor is composed of spaces (such as rooms and corridors), where each space contains features such as furniture, windows, and so on. Figure 9.15 illustrates the structural and semantic hierarchical models for a section of the Gulbenkian Museum. Semantic containment is based on the logical grouping of building components regardless of the physical relationship between those spaces. For example, two different art gallery rooms may be located

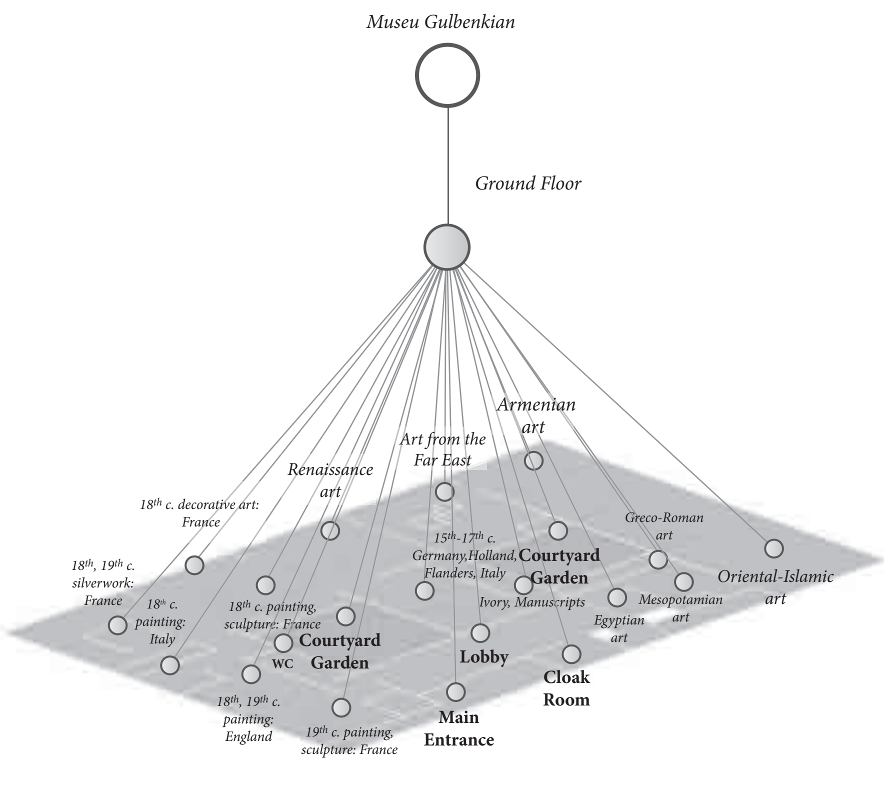

(a) Structural Hierarchy

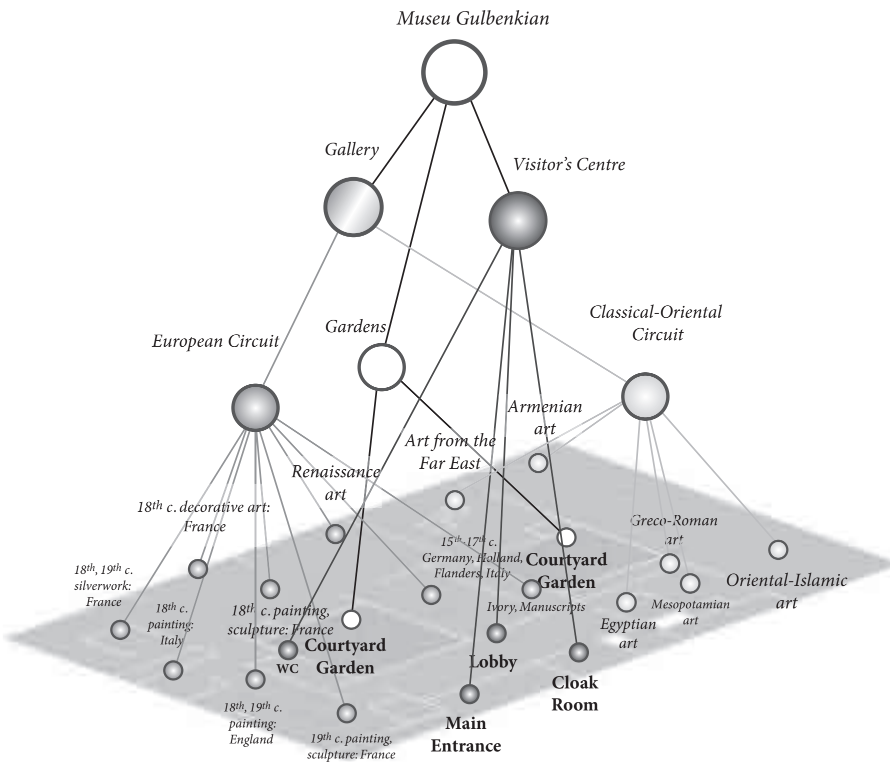

(b) Semantic Hierarchy

FIGURE 9.15 Hierarchical models of the Gulbenkian Museum.

on separate floors (or in completely different sections of the building), but are logically grouped together as belonging to the same art era. Another example is the set of components that are associated with *utilities* at the highest abstraction level such as water, heating, electricity, or the ventilation system, each of which, in turn, consists of utility-specific components distributed throughout the building such as air vents, ventilation shafts, terminals for controlling airflow, and so on.

9.5 Spatial computing within assistance systems

In Section 3, we presented the spatio-linguistic conceptualizations that occur within the range of the select application domains, namely, the architecture and media pre-production design domains, and the real-time emergency assistance scenario. Section 9.4 introduced the key notion of the 'structural form' that accrues as an abstraction mechanism for the formalism of spatial structure. Here, the focus was on the multi-perspective, qualitative, and multi-modal characterizations that constitute one model for our notion of structural form. In this section, we now turn to some instances of the types of spatial reasoning tasks that may be achieved with the abstraction and formalization mechanisms illustrated so far: Section 9.5.1 attempts to take stock of the paradigmatic underpinnings and fundamental problems within the scope of spatial reasoning. Sections 9.5.2–9.5.4 provide concrete examples grounded to the application domains introduced in the chapter.

9.5.1 Spatial computing: guiding principles and fundamental problems

The kinds of fundamental reasoning tasks that may be identified within the purview of spatial reasoning span a wide spectrum, for example including reasoning patterns such as spatial property projection, spatial simulation, spatial planning (e.g. for configuration problems), explanation with spatial information (i.e. causal explanation, hypothetical reasoning) to name a few. Both within and beyond the range of domains identified in this chapter, these are reasoning problems that involve an inherent interaction between space, actions, events, and spatial change in the backdrop of domain-specific knowledge and commonsense knowledge about the world (Bhatt, 2010).

Our notion 'spatial computing' has at least two semantic interpretations: (1) it refers to computing spatial relations, and (2) it refers to using spatial structures to do the computing. When we use the notion 'spatial computing', we actually refer to both meanings simultaneously, that is, we use spatial structures to compute spatial relations. For instance, a simple and well-known example of spatial computing heavily used in architecture is constructive geometry: here we use a flat sheet of paper, an architect's plan, as a spatial structure that is made to correspond in certain aspects to the spatial structure of the floor of a building. In particular, angles on the plan are identical to the corresponding angles of the building. Distances typically intentionally are not chosen to be identical, to make it easier to handle the plan and to obtain a good overview of the depicted structure; however, the ratios of corresponding distances and the ratios of corresponding areas are identical.

The important aspect of the architect's representation of the floor layout of the building is that many correspondences 'automatically' are generated by their plan without requiring any computing time: in the moment the architect has correctly

[15] A rather disconnected, but interesting analogue may be identified here in the context of the ISO 216 Standardisation of paper sizes: all ISO paper sizes in the *A*, *B*, and *C* series have the same aspect ratio of $1{:}\sqrt{2}$. If a sheet with this ratio is divided into two equal halves parallel to its shortest sides, then the halves will again have the same ratio. The practical advantages of this are many, but most importantly, in automatic scaling (e.g. *A*3-to-*A*4 on copiers) without compromising layout, paper, or space wastage, and estimation (e.g. postal and publisher estimates, organization and binding ease for libraries), etc. All this is achieved by just one structural constraint, which in some sense, is similar to our semantic characterization of spatial structure being a rather powerful computer in itself.

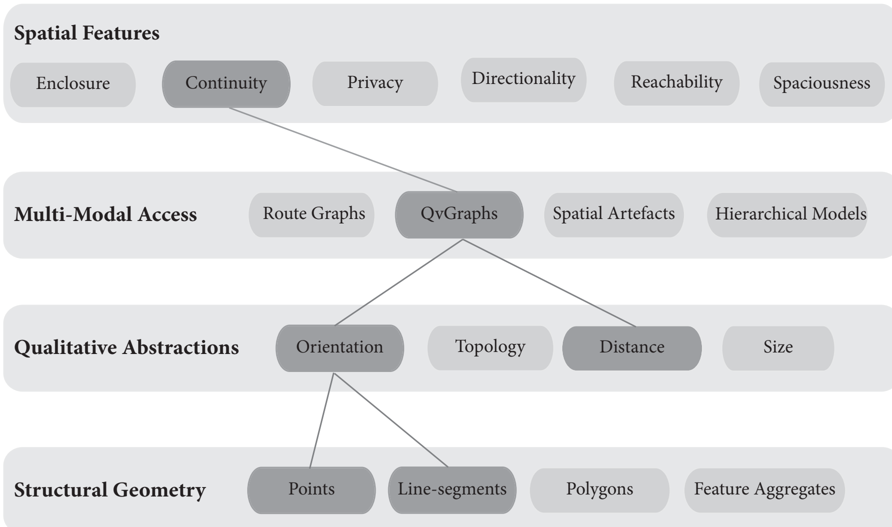

FIGURE 9.16 Spatial Qualities within the Architecture Design Assistance Domain as Viewed Through Multi-Modal Lenses

drawn the lines for parts of the layout, the corresponding angles, the corresponding ratios between line segments, the corresponding areas, and all other spatial correspondences are established. In other words, the architect's plan is an extremely powerful 'computer'.[15] The computing power of the architect's plan comes at a cost, though: you can't automatically do all the things that a regular computer can do; it is specialized at computing spatial relations.

9.5.2. *The behaviour and function in spatial structure, as viewed through multi-modal lenses*

The illustration in Figure 9.16 is an elaboration of the concept of multi-perspective semantics from Section 9.4.2: the top-most component comprises spatial features that may be identified at an initial design conception stage. For instance, it may be desired that a certain (set of) spatial structure(s) or its components may be desired to fulfil one or more of the qualities exemplified in Figure 9.16.

As one goes lower down the abstraction hierarchy of Figure 9.16, one moves close to the quantitative perspective of the design tool, which is where the precise geometry of a spatial structure resides. The basic idea is short, and simple: high-level spatial features and qualities that exist within the conceptualization of a human (e.g. designer, creative artist) correspond to one or more spatial structures at a lower level of abstraction. Some examples from the domain of architecture follow.

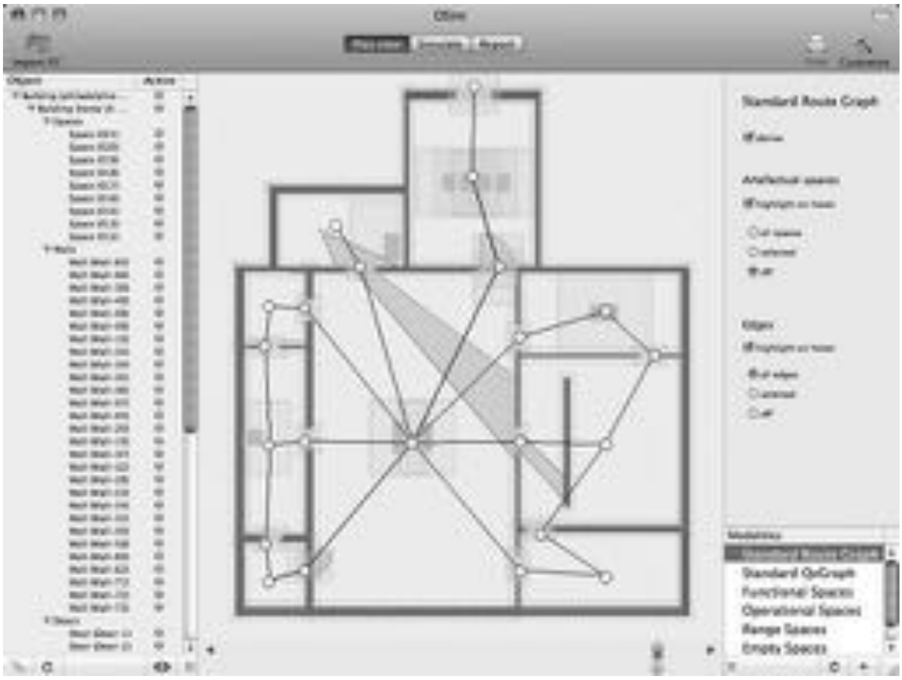

FIGURE 9.17 DSim – A Design Assistance System (prototype; 01.01.2011)

9.5.2.1 Privacy (Security) A typical design requirement may entail that certain parts of the environment may or may not be visible or readily accessible. For instance, it may be desired that the 'WashRoom' is as isolated as possible from other work areas, and that it may not be within the reach of sensing apparatus such as in-house cameras. As an example, such a constraint may be directly encoded at a higher level of abstraction within a rule-based programming mode (also see Figure 9.17):[16]

secure by (Door, Sensor):-
structural_geometry(Door, SGeom),
operational_space(SGeom, OpSpace),
range_space(Sensor, RgSpace),
topology(OpSpace, RgSpace, inside).

The example constraint utilizes the spatial abstractions that were introduced in Section 9.4.3 by way of the *spatial artefacts*: that is, the requirement of visibility or

[16] Indeed, such a rule-based model would technically rely on the underlying representational framework that has been discussed herein. Precise details are not relevant for this chapter, but it may be noted that we realize such rule specifications within the Constraint Logic Programming framework (Bhatt et al., 2011). Also refer to the concluding discussions on this topic in Section 11.6.

invisibility is formulated in terms of a topological constraint between two spatial artefacts, namely, the *range space* (of a sensor), and the *operational space* (of the door). Note the manner in which even a simple rule such as this utilizes concepts and data structures from three different perspectives: the quantitative geometry coming from the CAAD model, the qualitative abstractions by way of spatial artefacts and topological relationships (e.g. *inside*), and the conceptual level of the designer, consisting of elements such as 'Doors' and 'Sensors'.

9.5.2.2 Continuity Continuity among a set of entities or locations may be, for instance, identified as mutual visibility among the entities and locations under consideration. In addition to mutual-visibility as an interpretation of continuity, one may expect additional constraints involving relative and absolute positional constraints among the entities. Additionally, one may further refine the notion by the inclusion of distance constraints, for example '*X should not only be visible from Y, it should also not be too far away*...'. At a lower level of abstraction, the notion of continuity therefore translates to a set of visibility, orientational, and distance constraints over a spatial structure, which is precisely the perspective offered by the modality of a QvGraph (see Figure 9.11(a); Section 9.4.3).

9.5.2.3 Spaciousness One may interpret a high-level spatial feature such as *spaciousness* as denoting the sense of volume and openness felt by an occupant, that is, spaciousness too may be interpreted as being related to the notions of visibility and the arrangement of objects within a room (Flynn et al., 1973; Flynn, 1977). For instance a spacious environment can have many objects around the perimeter of the room, but critically must have no large objects, or very few, in the central region of the room. Centrally located objects can occlude mutual visibility of large portions of a room, and objects positioned on walking paths can create a sense of clutter (Key, 2009).

Our objective in the above examples pertaining to spatial features has been to illustrate the usability of our multi-perspective and multi-modal abstractions as an interface between structural form, and the behaviour and function that it entails. The detailed examples for spatial features such as continuity and spaciousness are rather involved from an implementation viewpoint and have therefore been omitted for the ongoing discussion.

9.5.3. Expected navigation and movement patterns through art exhibition spaces

The manner in which a museum space is navigated by a visitor will have a direct impact on their experience of the exhibits (Wineman and Peponis, 2010). Moreover, structural features of the layout often dictate, to a greater or lesser degree depending on the agenda of the visitor, how the space is explored, and in turn, affect the degree of exposure that particular exhibition pieces are afforded. For example, Melton

(1933) presented these ideas in a seminal article after conducting a number of studies on movement patterns of museum visitors.

Consider the situation where an art director of the Gulbenkian Museum is planning to hold a temporary exhibit. The art director needs to make decisions about the placement of the temporary pieces to get the desired effect; for example the art director may want to

- elicit a sense of impact and boldness with the new pieces, or alternatively, introduce them in a subtle way;
- maximize the exposure of temporary exhibits, or alternatively, strike a balance between the prominence of the new exhibits and the permanent works.

The SAS with spatial modalities can be used to predict the expected movement patterns of visitors and thus assist in the art director's task of placing the new exhibits to evoke the desired impression in the visitors. Modelling the principles of movement patterns presented by the architecture research community requires combinations of spatial modalities.

At a global level, as people move between different spaces within the museum they tend to establish a primary set of paths referred to collectively as the *skeleton* (Kuipers et al., 2003) of the environment in their comprehension of the space. Using this skeleton as a reference, visitors explore more localized collections of gallery rooms. Kuipers et al. (2003) propose that the skeleton that emerges as a person explores an environment tends to consist of *major* paths and locations that have a relatively high degree of topological connectivity to other paths and locations (i.e. boundary relations). Wineman and Peponis (2010) link the notion of a skeleton to research on the movement of visitors in museums such as Choi (1999); they emphasize the role of *accessibility* and *visibility* between spaces in characterizing major, integrated paths that form the skeleton (Wineman and Peponis, 2010). By analysing combinations of spatial modalities, in particular, route graphs, QvGraphs, and semantic hierarchical models, an architect can build up a picture of the accessibility and visibility afforded by the museum layout. Thus, the architect can make reasonable predictions about visitor movement patterns, for example, by automatically deriving a skeleton according to connectivity and visibility metrics. Moreover, researchers have investigated structural characteristics that influence the relative duration of occupancy of a space; for example, Choi's (1999) studies lead to the following principle (as expressed succinctly in Wineman and Peponis (2010):

Principle 1: Visitors stop more often in spaces that have greater visual connections to other spaces; they also stop in spaces that are visually connected to [spaces on the skeleton]

Spaces in which visitors are likely to stay for relatively longer periods can be identified by analysing route graphs and derived skeletons in conjunction with

QvGraphs. Thus, the art director can use the spatial modalities to determine the expected global movement patterns of visitors through the museum, and then decide where to place the new exhibits according to their desired aims. For example, by placing new exhibits in areas of high visitor traffic, the art director can ensure high exposure and conjure a sense of impact surrounding the new pieces.

At a local level, people move through a gallery room based on the layout of exhibits and the shape of the gallery room. Spatial assistance systems can employ key spatial modalities to model these movement patterns. We will consider a number of principles that have been proposed in the architecture research community for visitor movement behaviour and model these using as a vector flow graph. Bitgood (1995) has compiled a collection of principles that govern visitor movement within a gallery room based on well-known architecture research. For example:

Principle 2: Visitors tend to turn in the direction of the closest visible exhibit, all other factors being equal [. . .]

This indicates that *distance* plays a role in exhibit selection. This can be used to divide exhibits into those that are either 'near' or 'far' from a given visitor location based on a threshold shortest-path distance through the route graph, where 'near' exhibits have priority over 'far' exhibits. Exhibits that are 'near' are considered to be in the same equivalence class with respect to distance, and thus the following principles can then take effect in determining further selection.

Wineman and Peponis (2010) have proposed a collection of principles that govern visitor movement in very open-plan gallery spaces (i.e. where the visitor is not explicitly guided along a restricted set of paths):

Principle 3: the more accessible an exhibit element is from all other exhibit elements, the more likely it is to be visited. This provides the [. . .] hierarchy of the likelihood that an exhibit will be perceived in spatially guided movement.

Accessibility refers to how *easy* it is to navigate to an exhibit from other locations in the room. Using the route graph where each node corresponds to an exhibit, accessibility of a given exhibit can be measured by the graph-theoretic centrality value (many possible centrality metrics can be experimented with, for example the degree, betweenness, and closeness). Thus, if a visitor is presented with a selection of exhibits, all other factors being equal, they will tend to choose the most accessible exhibit.

Bitgood (1992) has proposed a precedence ordering on the decisions that museum visitors make, again based on a number of well-known architecture studies (typically taking place in North American or Western European museums).

Principle 4: [. . .] visitors exit a gallery by the first open door they encounter

If a visitor is near an exit door, then they will take it.

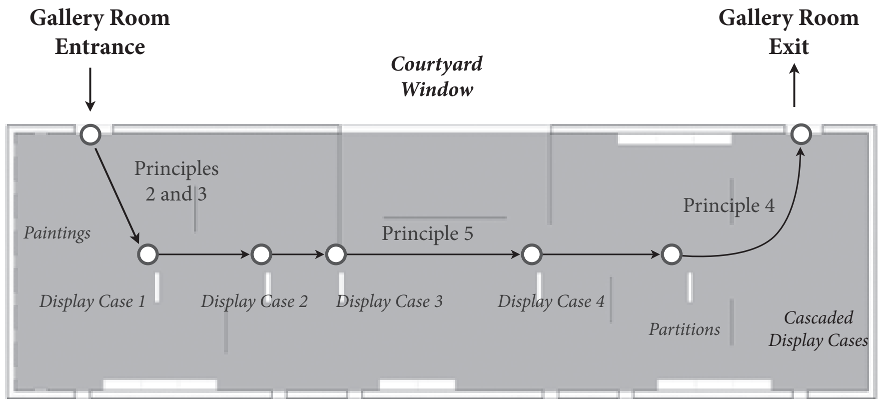

FIGURE 9.18 Expected visitor movement pattern through the Oriental-Islamic and Armenian gallery rooms.

Principle 5: [. . .] people tend to walk in the same direction. Thus, if a visitor enters a gallery along the left-hand wall, [they] continue walking along this wall

This introduces a notion of visitor trajectories, such that the direction of the path defined by the previously visited exhibit e_1 and the current exhibit e_2 will influence the decision for the next exhibit e_3.

Figure 9.18 illustrates the application of these principles in generating a flow vector graph of the expected visitor movement pattern. We can observe that the central display cases dominate the paintings on the right-hand side of the entrance as visitors enter the gallery room due to proximity and accessibility (Principles 2 and 3). As the visitor travels down the gallery room, they are expected to exclusively take the left-hand path; this is because a movement trajectory is established (Principle 5) by the regularly placed display cabinets (each providing a natural *next location* due to proximity and accessibility, i.e. Principles 2 and 3). As soon as the exit becomes visible the visitors leave the room (Principle 4), ignoring the cascading display cabinets along the back wall. Using this predicted movement pattern, the art director can decide whether to place new temporary exhibits along the expected path, or perhaps position them in a more subtle position such as the courtyard window, or near the paintings at the entrance.

9.5.4. Scenario and narrative completion by spatio-temporal abduction

Re-consider the illustration in Figure 11.6 for the domain of automatic cinematography that was introduced in Section 9.3.2: the world consists of three point-abstracted entities—2 *avatars* and 1 virtual *camera*.[17] Suppose that container space

[17] The third entity in the simulation is a virtual camera that records the other two entities in the scene, and hence is not visible within the three-dimensional illustration of Figure 9.6.

is modelled as a discrete grid world together with relative orientation relationships among the entities as per the partitioning scheme of the Single-Cross Calculus (Freska, 1992). For this discussion, further suppose that the *camera* is the only entity that is able to move, that is, change location from one grid-cell to another.

For a scenario such as this, spatio-temporal abduction serves as a basis of *scenario and narrative completion*, and for this particular example, the derivation of ideal *camera placements* serves as a side-effect of the abduction process. Figure 9.19 consists of a *narrative* (completion) from time points t_1 to t_{12}, denoting an *abduced* evolution of the system, as represented by the sequence of qualitative state descriptions for *two* stationery and *one* moving entity. For clarity, images from a three-dimensional simulation are included together with the relational illustrations for each of the time points. From an initial narrative description consisting of information about only some of the time points,[18] the narrative completion has been abduced on the basis of available *camera actions—pan*, *zoom*, *move*—and pre-specified knowledge or heuristics, referred to as *film idioms*, about desired camera placements, for example *establishing shot, external shot, mid-shot, close-up*, and so forth. In this example, the resulting narrative may be used by a virtual reality or an automatic cinematography system to generate automatic visualizations for a script.

With respect to the automatic cinematography example, it is easy to intuitively infer the general structure of causal explanation (by abduction) within spatial information. Consider the illustration in Figure 9.20 for a hypothetical (e.g. branching) situation space that characterizes the complete evolution of a system. In Figure 9.20—the situation-based history $< s_0, s_1, \ldots, s_n >$ represents one path, corresponding to a actual time line $< t_0, t_1, \ldots, t_n >$, within the overall branching-tree structured situation space. Given incomplete narrative descriptions, for example corresponding to only some ordered time points (such as in Figure 9.19) in terms of high-level spatial (e.g. topological, orientation) and occurrence information, the objective of causal explanation is to derive one or more paths from the branching situation space, that could best-fit the available narrative information. Of course, the completions that bridge the narrative by interpolating the missing spatial and action/event information have to be consistent with domain-specific and domain-independent rules (e.g. pertaining to dynamics).

A formalization of the above stated problem of scenario and narrative completion by spatio-temporal abduction in the Event Calculus may be consulted in Bhatt and Flanagan (2010). The motivations and broad research questions underlying the approach may be referred to in Bhatt (2010).

[18] These are, for instance, (implicitly) available from linguistic descriptions about *acts* and *scenes* within a drama or film script. Here, the progression of the script can be thought of as an imaginary evolution of the system.

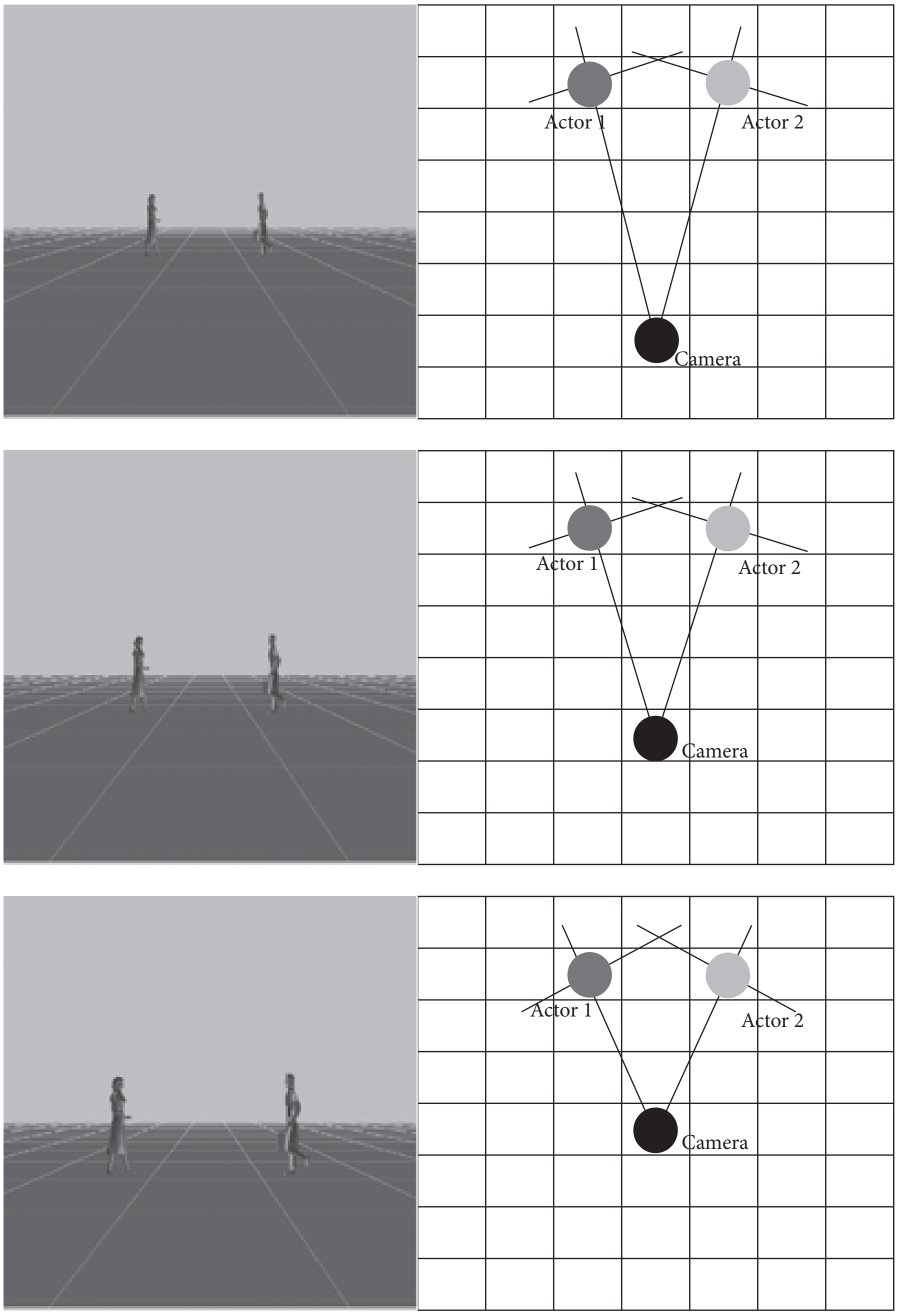

FIGURE 9.19 Scenario and Narrative Completion by Abduction. Source [7].

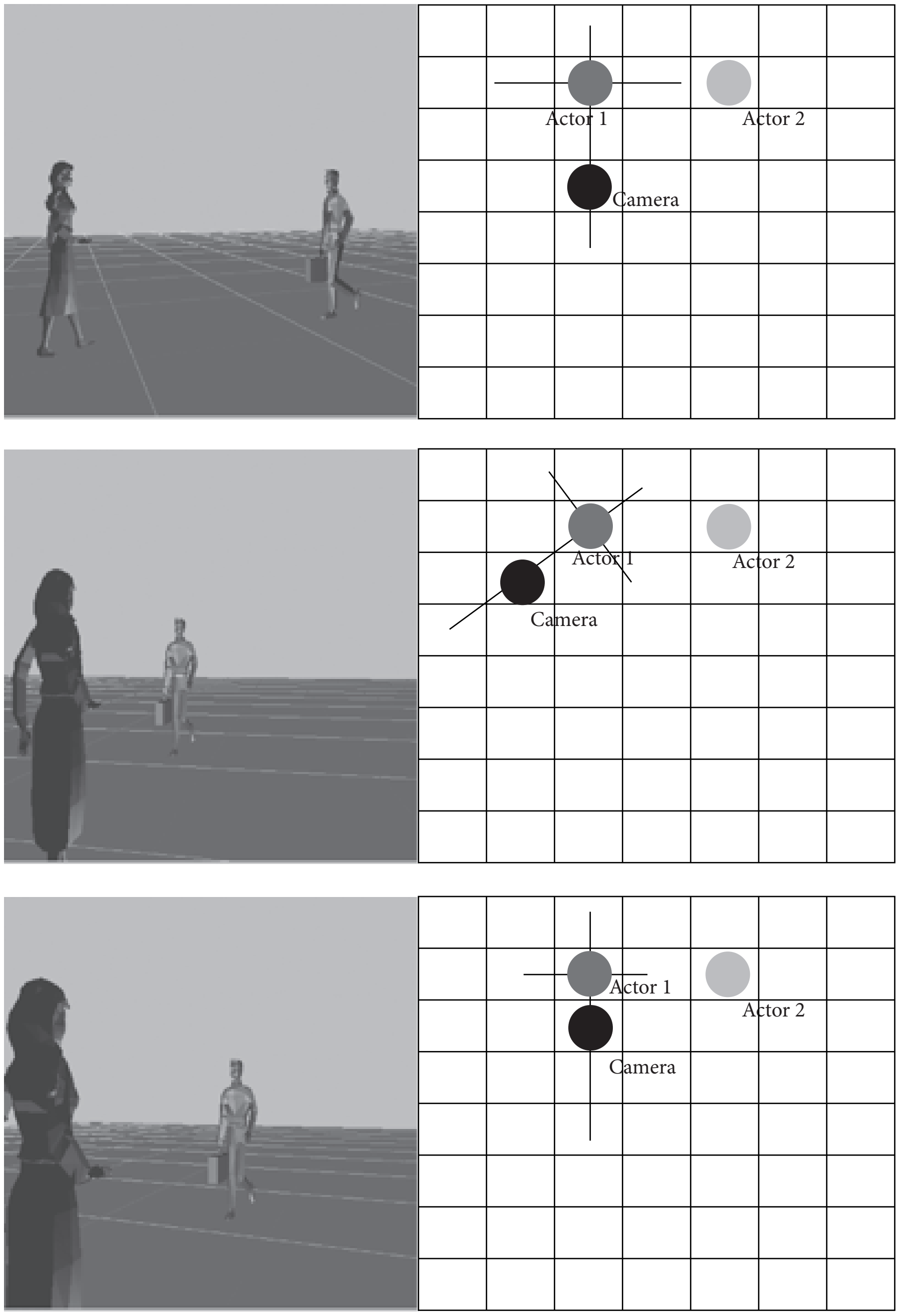

FIGURE 9.19 Continued

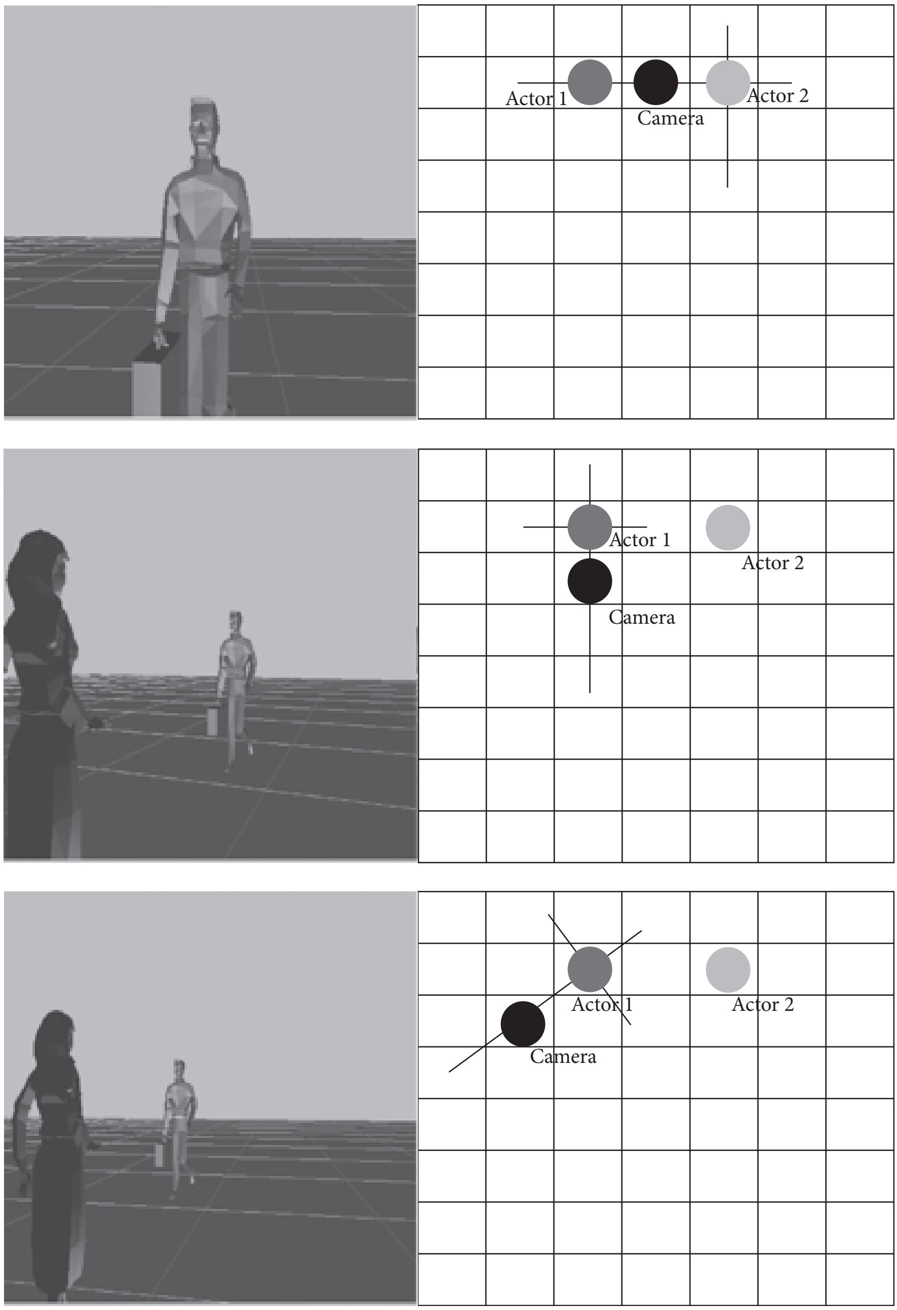

FIGURE 9.19 Continued

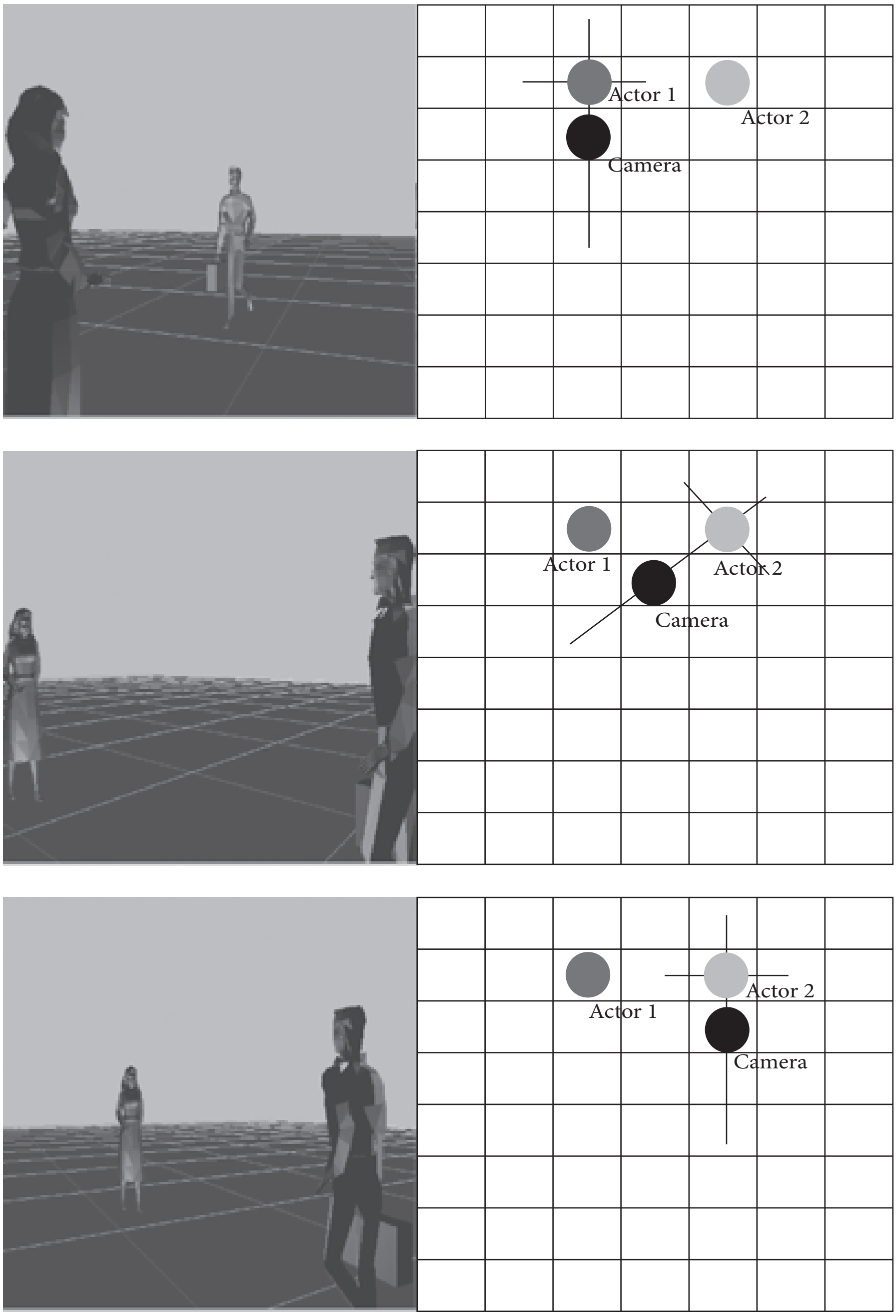

Figure 9.19 Continued

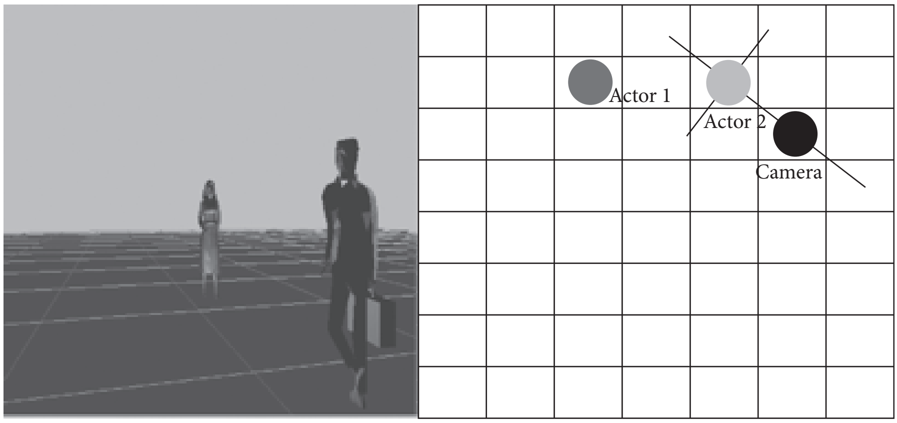

FIGURE 9.19 Continued

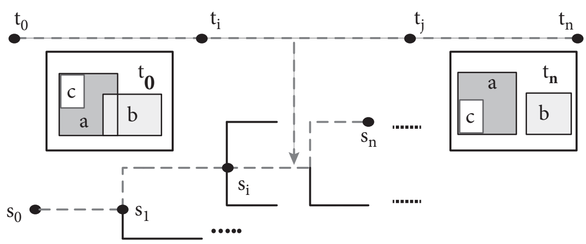

FIGURE 9.20 Branching / Hypothetical Situation Space

9.6 Summary and discussion of chapter

Spatial assistance systems aim to transfer the cognitive stress involved in a human analytical activity onto a system, by externalizing and operationalizing the decision-making processes involved therein. From the viewpoint of human *assistance*, *assurance*, and *empowerment*, this chapter considered a range of application domains, including architecture design assistance, creative media assistance, and navigation assistance (e.g. emergency scenarios). This chapter focused on the spatial informatics concerned with the conceptual, formal, and computational aspects of 'space' within this range of *spatial assistance systems*. Specifically, the main contributions of this chapter are:

1) identification of the nature of human-centred spatio-linguistic conceptions that occur within a specific class of application domains;
2) development of a formal framework that may be used as an interfacing mechanism between the spatio-linguistic conceptualization of humans and its qualitative abstraction as per our interpretation of structural form;

3) enabling fundamental reasoning patterns, as identifiable with respect to a particular philosophy of general spatial reasoning within systems that creatively or constructively, assist, assure, and empower humans.

Whereas the above were illustrated in the context of diverse application scenarios, the discussions focused on the domain of architecture design assistance in particular. This is because spatial design as a problem-solving activity involves rich, real-world problem characterizations and solutions: the domain also enables concrete specification and illustration of problems and their solutions. The other two domains, involving creative media design, and real-time emergency assistance were chosen because in addition to being practically relevant, they offered significant breadth in order to illustrate the generality and transferability of the basic ideas concerning structural forms, multi-perspective, and multi-modal abstractions, and the computational problems addressed in the chapter.

9.6.1 Further pointers to literature

Much has been left out of the chapter: a precise definition and formalization of structural form—each constituent component thereof—and a formal treatment of the computational aspects concerning checking for functional design consistency using rule-based specifications has been omitted. For example, rule-based specifications, which we have implemented within the Constraint Logic Programming (CLP) framework (Jaffar and Maher, 1994), present their own set of fundamental challenges involving the realization of a 'spatial semantics' such that rules may refer to spatial entities and relationships (Bhatt et al., 2011). Similarly, an elaboration of the rather contrived notion of (logic-based) spatio-temporal abduction has not been included. Spatio-temporal abduction, which can be implemented (Bhatt and Flanagan, 2010) using the Event Calculus formalism (Kowalski and Sergot, 1986), has only been described with respect to the computational structure that it acquires; much more may be elaborated on in this regard, for example, with respect to its precise formalization, the integration of the semantics of a qualitative spatial calculus within the event calculus, and the general applicability of spatio-temporal abduction in domains outside of the one discussed here (e.g. activity abduction in smart environments, geospatial dynamics). For the interested reader these aspects, and the general agenda that underlies and inspires this line of research may be consulted in Bhatt (2010), Bhatt and Freska (2010), and Bhatt and Loke (2008)

Finally, although we presented the concept of multi-perspective semantics and multiple modalities in sufficient detail, the industrial underpinnings and relevance of this work were not presented in the chapter: at a practical level, we have developed a multi-perspective, multi-modal spatial data access framework (Schultz and Bhatt, 2010) designed to serve the informational and computational requirements of architectural design assistance systems that are intended to provide intelligent spatial

OUP UNCORRECTED PROOF – Page Proof, 11/6/2013, SPi

decision-support and analytical capabilities. In this context, we ensure interoperability with commercial tools by utilizing the stipulations of the Building Information Model (BIM) (Eastman et al., 2008) and the Industry Foundation Classes (IFC) (Froese et al., 1999). Exemplary prototypes are also in progress as a part of the design assistance tool *DSim* (Bhatt et al., 2010) that is being designed to offer reasoning, simulation, and visualization capabilities.