Int J Software Informatics, Volume 9, Issue 2 (2015), pp. 279–288
International Journal of Software and Informatics, ISSN 1673-7288

E-mail: ijsi@iscas.ac.cn
http://www.ijsi.org
Tel: +86-10-62661040

Computational Problem Solving in Spatial Substrates
— A Cognitive Systems Engineering Approach

Christian Freksa

(Bremen Spatial Cognition Center, University of Bremen, Germany)

Abstract The ability to perform spatial tasks is crucial for everyday life and of great importance to cognitive agents such as humans, animals, and autonomous robots. A common artificial intelligence approach to accomplish spatial tasks is to represent spatial configurations and tasks in form of detailed knowledge about various aspects of space and time. Suitable algorithms then use the representations to compute solutions to spatial problems. In comparison, natural embodied and situated agents often solve spatial tasks without detailed knowledge about geometric, topological, or mechanical laws; they directly relate actions to effects that are due to spatio-temporal affordances in their bodies and environments. Accordingly, we propose a paradigm that makes the spatio-temporal substrate an integral part of the engine that drives spatial problem solving. We argue that spatial and temporal structures in body and environment can substantially support (and even replace) reasoning effort in computational processes: physical manipulation and perception in spatial environments substitute formal computation. While the approach is well known – for example, we employ diagrams as spatial substrate for geometric problem solving and maps for wayfinding – the underlying principle has not been systematically investigated or formally analyzed as a paradigm of cognitive processing. Topology, distance, and orientation constraints are all integrated and interdependent in truly 2- or 3-dimensional space. Exploiting this fact may not only help overcome the need for acquiring detailed knowledge about the interrelationships between different aspects of space; it also can point to a way of avoiding exploding computational complexity that occurs when we deal with these aspects of space in complex real-world scenarios. Our approach employs affordance-based object-level problem solving to complement knowledge-level formal approaches. We will assess strengths and weaknesses of the new cognitive systems paradigm.

Key words: spatial substrates; cognitive system architecture; embodied and situated cognition; knowledge in the world; perception and action in space

Freksa C. Computational problem solving in spatial substrates – a cognitive systems engineering approach. *Int J Software Informatics*, **Vol.9, No.2 (2015): 279–288. http://www.ijsi.org/1673-7288/9/i218.htm**

Cordially dedicated to Bernd Krieg-Brückner.

Corresponding author: Christian Freksa, Email: freksa@uni-bremen.de
Received 2015-01-02; Revised 2015-05-14; Accepted 2015-05-20.

1 Introduction

This work focuses on a special class of real-world problems that is of particular significance for cognitive agents such as humans, animals, and autonomous robots: spatial and temporal problems in physical environments. Spatio-temporal problems share basic structural properties that have been intensively studied over the past twenty years[12] and are quite well understood today on the information processing level. Despite this similarity, the best solutions to different types of spatio-temporal problems employ a considerable variety of approaches and tools. These tools include abstract computational algorithms to solve geometric problems and concrete physical tools like screwdrivers and pliers that may serve to exploit spatial affordances.

Some of the best approaches for human spatial problem solving make heavy use of the physical object level rather than solving problems entirely on the abstract information level. For example, when a cognitive agent instantiates the route instruction *turn left at the next intersection*, he, she, or it does not require a detailed mental representation of the environment; it also does not need to know and specify a precise turning angle before being able to follow the instruction. Rather, the instruction provides a coarse guide line and permits the agent to move in an unfamiliar environment by perception-based route following and to select a 'left' route from several alternatives that it may perceive in the vicinity of the intersection. The information about the turning angle is implicitly present in the spatial configuration that consists of the pose of the agent in relation to the route; the instruction can be followed by means of a short perception-action loop; this does not require that the turning angle ever be made explicit in a cognitive representation.

In other situations cognitive agents may prefer to have detailed spatial knowledge before starting a spatial action as it may be easier to solve the problem by reasoning than by spatial interaction. For example, when I lost my keys that I last used during a trip some while ago, it may be worthwhile to reconstruct the preceding sequence of events on the trip mentally; directly exploring the environment perceptually might also work, but it could be quite difficult or laborious, in the particular situation. Current robotic approaches predominately rely on detailed knowledge about environments and their properties represented in computer memory [e.g. Ref. [35]]. Embodied and situated cognitive agents are capable of operating both, on the information processing and on the physical object level and of combining both levels in smart ways. Perception and action operations serve as interfaces between the two levels; a memory serves to make information about the environment available to information processing in the absence of perceptual information and to store results of information processing for carrying out actions.

The classical model of cognition as an information processing activity that takes place entirely in the brain (respectively computer) is only one way of performing cognitive tasks[36]; it presupposes that a real-world problem has been comprehensively abstracted into a pure information-processing task. This assumption may be appropriate for routine tasks for which all necessary information is provided and which can be performed according to pre-existing standard patterns; for novel problems, however, a considerable part of the problem solving effort goes into ***finding the most appropriate approach, suitable tools, and identifying the information needed to solve the problems***; in many cases, a quite specific approach may be more appropriate and more efficient than a general approach. Specific approaches take into account particular features of the problem domain to a larger extent than general approaches that tend to abstract from specific characteristics. In general, cognitive agents have a considerable variety of approaches and tools they can use to attempt solving a given spatial problem. In real-world cognitive problem solving we are frequently confronted with problems where the identification of a suitable approach is a far more difficult task than the computation of a solution on the basis of a given approach; once an appropriate approach has been selected, the problem solving procedure itself may be straightforward.

Thus, we need to address the question of *how to find a suitable approach* to solve a given spatial problem. In general, finding a suitable approach to solve a novel problem is one of the most interesting and challenging problems for cognitive agents. For the specific domain of spatio-temporal problems, we have reasons to believe that the time is ripe to tackle this challenge; our belief is rooted in the fact that today we have a much better understanding of the properties of spatial and temporal relations and structures than twenty years ago. There is even hope that once the spatio-temporal reasoning challenges are tackled, we will be able to use the resulting approaches for dealing with non-spatial problems as well; this hope is rooted in the insight that human cognitive agents understand many problems through analogies[17] and metaphors[22]; thus, non-spatial problems may be solved by mapping them onto spatially constrained structures which may be easier to solve; this would be in contrast to generalizing spatial approaches to unconstrained domains where we would employ highly general approaches.

The spatial (and to a lesser extent the temporal) domain is particularly well accessible to autonomous mobile agents with visual, haptic, and auditory perception and memory as well as with moving, turning, and grasping capabilities. These capabilities enable the agents to flexibly interact with their environments; specifically, they can modify the parts and aspects of the environment they perceive and they can modify spatial configurations in the environment through their actions. In robotics, sensory capabilities have been successfully employed to avoid obstacles in similar ways as animals by implementing reflexes that do not require representations in the brain[7,33]. These capabilities have not been systematically investigated and exploited for cognitive systems architectures beyond obstacle avoidance, so far. In our research, we develop proof of concept implementations and demonstrations for solving spatio-temporal problems strategically by making use of spatio-temporal affordances.

A main research hypothesis for studying physical operations and processes in spatial and temporal form in comparison to formal or computational structures is that **spatial and temporal structures *in the body and the environment* can substantially support (and even replace) reasoning effort in computational processes**. A major observation we can make when we compare the use of different forms of representation is that the processing structures of problem solving processes differ see Ref. [23] and facilitate different processing mechanisms[32]. Structures that resemble the problem domain may result in a lower complexity class than structurally deviating representations as they can make direct use of the structural properties without a need for describing them[25].

A main objective of our work is to explore the scope of application of this principle. This will involve a representation-theoretic assessment of representational equivalence and similarity, both on the level of result and process equivalence (or similarity). We develop a framework to relate physical actions and perception activities to information processing activities, in order to assess the trade-off between physical and mental operations. Such a framework has long been missing in the debate surrounding diagrammatic vs. analytic reasoning. Our approach builds on well-established paradigms from cognitive science (e.g. 'knowledge representation theory'[28], 'affordances'[18], 'knowledge in the world'[27], 'conceptual neighborhood'[10]) and on research carried out in the collaborative research center SFB/TR 8 Spatial Cognition at the University of Bremen over the past twelve years.

2 Background and Motivation

AI research initially was concerned exclusively with mental aspects of cognitive systems, specifically with operations and processes that take place in the brain (respectively computer)[9]. Advances in robotics and knowledge representation have extended the scope of AI research to model perception and action processes, the (physical) bodies of agents, and the agents' spatial environments [e.g. Ref. [8]]. The rather general structures of abstract formalisms used for knowledge representation in computers allow describing arbitrary aspects of bodies and environments in detail and to reason about them, including spatial and temporal aspects.

While abstract reasoning about the world can be considered the most advanced level of cognitive ability, this ability requires a comprehensive understanding of mechanisms responsible for the behavior of bodies and environments. But many natural cognitive agents (including adults, children, and animals) lack a detailed understanding of their environments and still are able to interact with them rather intelligently. For example, they may be able to open and close doors in a goal-directed fashion without understanding the mechanisms of doors or locks on a functional level. This suggests that knowledge-based reasoning may not be the only way to implementing problem solving in cognitive systems.

In fact, alternative models of perceiving and moving goal-oriented autonomous systems have been proposed in biocybernetics and AI research to model aspects of cognitive agents e.g. Refs. [4,5,29]. These models implement perceptual and cognitive mechanisms that follow physical laws rather than formal representations that follow the laws of logics. Such systems are capable of reacting to their environments intelligently without encoding knowledge about the mechanisms behind the actions and without the associated computational cost.

In our spatial cognition research we have investigated the potential of *qualitative spatial relations*, of structure-preserving *schematic maps*, and of the role of intrinsically spatial structures for spatial reasoning and spatial problem solving[11,13,31]. A main result of this work is that structure-preserving representations can make direct use of spatial relations (e.g. spatial neighborhood, conceptual neighborhood, spatial order, and spatial orientation); without structure-preservation, these relations would have to be derived through knowledge-based processes in more abstract formal representations. Thus, spatial calculi that exploit structure-preserving representations can avoid the necessity of performing certain computational derivations.

Spatial cognition research also has been concerned with issues of resolution and granularity, both on a physical and on a conceptual level[20,15,26,30]. In knowledge representation, we must deal with the issue of level of detail on which we represent objects and configurations in order to solve certain problems. The finer the level of representation, the more problems we will be able to solve, *in principle*. But this comes at a cost: the more details we have to deal with, the more computation we have to invest. The corresponding problem spaces often suffer from combinatorial explosion that prevents tractability. Cognitive processes frequently process information from coarse to fine rather than from fine to coarse. These processes are directly supported by physical and spatial properties of their environments. For example, in vision, *coarse* corresponds to *distant* and *fine* corresponds to *close-up*; the same sensor adapts its 'representation' of the world simply by physically moving towards an object or away from it.

The field of *diagrammatic reasoning*[2,6,19] is concerned with problem solving by means of diagrams, a special form of spatial representations. A key issue here is the comparison between formal and diagrammatic representations and reasoning processes for the same underlying problems. Of particular interest is the equivalence / similarity between the reasoning procedure operating on the corresponding formal structure and the reasoning procedure operating on the spatial structure. Process equivalence has been mainly studied in comparing different formal systems[1]. Comparing processes operating on physical spatial structures with processes operating on formal structures poses an interesting challenge, as we will require a reference framework that includes information processing and re-configuration of spatial configurations.

The background of this work has been discussed in more detail in Ref. [13].

3 Approach

In our present research, we go an important step beyond previous work and introduce a paradigm shift: we do not only aim at preserving spatial structure in representations, but we also make use of identity preservation; in other words, we represent spatial objects and configurations by themselves or by *physical spatial models* of themselves, rather than by abstract representations. This has a number of advantages: we can avoid loss of information due to early representational commitments and may be able to integrate several representations in the course of the problem solving process: we do not have to decide beforehand which aspects of the world to represent in a certain way and which aspects to abstract from; this can be decided partly during the problem solving process. During this process, additional contextual information may become available that can guide the choice of the specific representation to be used.

Perhaps more importantly, objects and configurations frequently are aggregated in a natural and meaningful way; for example, a chair may consist of a seat, several legs, and a back; if I move one component of a chair, I automatically (and simultaneously!) move the connected components and the entire chair, and vice versa. Thus, physical as well as conceptual relations can be manipulated by physical actions of a cognitive agent. This feature is not intrinsically available in abstract representations where computational derivations are required to infer such relations. Thus, manipulability of spatial structures is not merely a property of physical objects; it is an important feature of cognitive processing. Therefore we will include manipulation of spatial relations as a central component of our cognitive architecture.

Spatial manipulation is important for cognitive agents in at least two ways: (1) for active perception[3] and (2) for spatial problem solving. Active perception refers to attention processes and to perceptual response patterns in response to physical stimulation of the environment[23]. It also can be used to change an agent's egocentric spatial reference frame in order to obtain a more suitable perspective on a given spatial configuration for solving a problem. The following example presents the basic idea.

4. Example

4.1 Creating a suitable spatial reference frame

Suppose an agent's task is to determine visually (without a depth sensor) whether a tree is on its side of a fence or on the other side (Fig. 1(a)); a classical image analysis approach could use depth clues in the 2D projection of the 3D configuration to infer whether fence or tree is closer to the agent. Problem: the essential depth dimension is only weakly represented in this 2D projection. Spatial approach: Select a spatial reference frame that highlights the essential dimension; this is achieved by relocating the agent such that the essential dimension is projected prominently onto the image of the configuration (Fig. 1(b)); now the task can be solved by considering only one dimension on the image, as the previous depth dimension has been mapped to the perceptually better accessible width dimension by spatial transformation in the problem domain.

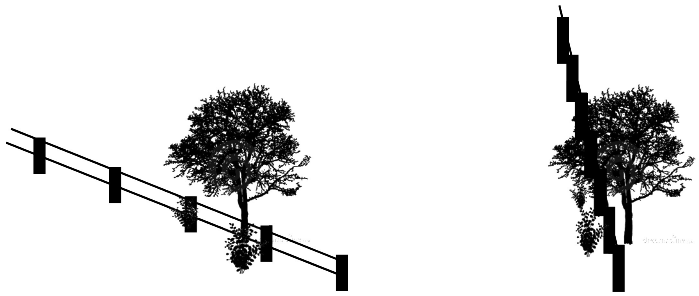

(a) (b)

Figure 1. (a) Hard visuo-spatial decision problem, (b) The same problem presented in a suitable spatial reference frame.

This specific example employs physical action only to modify perceptual acquisition of information from the environment without changing the 3-dimensional scene of interest. Note, however, that the 2-dimensional projection of the scene (that is typically the only information available for visual scene analysis) has considerably changed. Other kinds of manipulations would actually change the physical scene of interest in order to simplify or solve the spatial problem at hand.

We focus on spatial and spatio-temporal tasks that are directly accessible by perception and allow for manipulation by physical action. This is the domain we understand best in terms of computational structures; we have well established and universally accepted reference systems to describe and compute spatial and temporal relations. The limitation to spatial tasks may turn out less severe as it may seem initially: numerous non-spatial problems can be transformed into equivalent spatial problems where the spatial structure may support the problem solving process. Human problem solvers make use of problem spatialization for example when visualizing a linguistically specified problem in form of a diagram in order to better grasp the problem and/or to be better able to formalize it for formal problem solving. Depending on the spatial representation chosen for the diagram, it may be easier or harder to grasp or formalize the problem.

The main hypothesis of our approach is that the 'intelligence' of cognitive systems is grounded not only in specific abstract problem solving approaches, but also – and perhaps more importantly – in the capability of recognizing characteristic problem structures and of selecting particularly promising problem solving approaches for given tasks. Formal representations generally do not facilitate the recognition of such structures due to a bias inherent in the abstraction. This is where *mild abstraction* can help as it abstracts only from few aspects while preserving important structural properties.

The insight that spatial relations and physical operations are strongly connected to cognitive processing will lead to a different division of labor between the perceptual, the representational, the computational, and the locomotive parts of cognitive interaction than the one we have been pursuing in artificial intelligence: rather than putting all the 'intelligence' of the system into the computer, the proposed approach aims at putting more intelligence into the interactions between components and structures of a cognitive system as well as into the structure of the problem representation. More specifically, we aim at exploiting intrinsic structures of space and time to reduce the complexity of computation.

We argue that a flexible assignment of physical and computational resources for cognitive problem solving is closer to natural cognitive systems than the almost exclusively computational approach; for example, when we as cognitive agents search for certain objects in our environment, we have at least two different strategies at our disposal: we can represent the object in our mind and try to imagine and mentally reconstruct where it could or should be – this would correspond to the classical AI approach; or we can visually search for the object in our spatial environment. Which approach is better (or more promising) depends on a variety of factors including memory and physical effort required; frequently a clever combination of both approaches will be best.

We plan to develop and implement a proof of concept for the proposed approach to spatial problem solving through simulations of the perception and manipulation processes as well as through physical agent models, e.g. as generated by a 3D printer[14]. The research is primarily conceived as basic research in cognitive systems engineering: we want to identify and relate an inventory of cognitive principles and ways of combining them to obtain cognitive performance in spatio-temporal domains.

This project brings together perspectives from a variety of disciplines: (1) the cognitive systems perspective, which addresses the cognitive architecture and trade-offs between properties of physical structures and properties of their descriptions; (2) the formal perspective, which characterizes and analyzes the resulting structures and operations; (3) the engineering perspective, which constructs and explores varieties of cognitive system configurations; and (4) the psychological-empirical perspective, which relates the effects of different system behaviors to those of natural agents. In the long term, we see potential technical applications of physically supported cognitive configurations for example in the development of future *intelligent materials* (e.g. 'smart skin' where spatially distributed computation is required that needs to be minimized with respect to computation cycles and energy consumption and more robust and adaptable artificial agents, which can deal with unknown environments)[21,34].

For this project, we can build on extensive research on spatial and temporal relations, their representation in memory, and with qualitative spatial reasoning in the framework of international interdisciplinary spatial cognition research. Naturally, the proposed approach will not be as broadly applicable as some of the approaches we have pursued in classical AI research as it focuses on spatial and temporal structures; but the approach promises to discover broadly applicable cognitive engineering principles for the design of tomorrow's intelligent agents. Our philosophy is to understand and exploit pertinent features of space and time as modality-specific properties of cognitive systems that enable powerful specialized approaches in the specific domain of space and time. Since space and time are most basic for perception and action and ubiquitous in cognitive processing, we believe that understanding and utilizing their specific structures will be particularly beneficial.

Foundations of the approach have been outlined in more detail in Ref. [16]. A substantially extended description of this work with numerous examples has been published as Strong spatial cognition in the Conference on Spatial Information Theory 2015[37].

Acknowledgements

I thank anonymous reviewers for helpful critical comments and helpful pointers on an earlier version of the paper. Generous support by the German Research Foundation (DFG) for the SFB/TR 8 Spatial Cognition project R1-[ImageSpace] is gratefully acknowledged.

References

[1] Alves De Madeiros AK, Weijters AJMM. Process equivalence: Comparing two process models based on observed behavior. International Conference on Business Process Management. LNCS 4102. Berlin. Springer. 2006. 129–144.

[2] Anderson M, Furnas G. Relating two image-based diagrammatic reasoning architectures. In: Goel AK, Jamnik M, Narayanan NH, eds. Diagrammatic Representation and Inference. Berlin.

Springer. 2010. 128–143.
[3] Bajcsy R. Active perception, Proc. IEEE, 1988, 76: 996–1005.
[4] Braitenberg V. Vehicles: Experiments in synthetic psychology. Cambridge, MA: MIT Press, 1984.
[5] Brooks RA. Intelligence without representation. Artificial Intelligence, 1991, 47: 139–159.
[6] Chandrasekaran B. Multimodal cognitive architecture: Making perception more central to intelligent behavior. Proc AAAI. 2006. 1508–1512.
[7] Cruse H, Dautenhahn K, Schreiner H. Coactivation of leg reflexes in the stick insect. Biological Cybernetics, 1992, 67: 369–375.
[8] Davis E. Representations of Commonsense Knowledge. San Mateo, CA: Morgan Kaufmann, 1990.
[9] Feigenbaum E, Feldman J. Computers and Thought, New York: McGraw-Hill, 1963.
[10] Freksa C. Conceptual neighborhood and its role in temporal and spatial reasoning. In: Singh M, Travé-Massuyès L, eds. Decision Support Systems and Qualitative Reasoning, Amsterdam, North-Holland, 1991: 181–187.
[11] Freksa C. Qualitative spatial reasoning. In: Mark DM, Frank AU, eds, Cognitive and Linguistic Aspects of Geographic Space. Dordrecht. Kluwer. 1991. 361–372.
[12] Freksa C. Spatial and temporal structures in cognitive processes. In: Freksa C, Jantzen M, Valk R, eds. Foundations of Computer Science. Potential – Theory – Cognition. LNCS 1337. Springer. Berlin. 1997. 379–387.
[13] Freksa C. Spatial Computing – How spatial structures replace computational effort. In: Raubal M, Mark D, Frank A, eds. Cognitive and Linguistic Aspects of Geographic Space. Springer. Berlin. 2013. 23–42.
[14] Freksa C. Understanding cognition through synthesis and analysis. Int J Software Informatics, 2013, 7(1) : 3–18.
[15] Freksa C, Barkowsky T. On the relation between spatial concepts and geographic objects. In: Burrough P, Frank A, eds. Geographic Objects with Indeterminate Boundaries. London. Taylor and Francis. 1996. 109–121.
[16] Freksa C, Schultheis H. Three ways of using space. In: Montello DR, Grossner KE, Janelle DG, eds. Space in Mind: Concepts for Spatial Education. Cambridge, MA. MIT Press. 2014.
[17] Gentner D. Structure-mapping: A theoretical framework for analogy, Cognitive Science, 1983, 7(2): 155–170.
[18] Gibson JJ. The Theory of Affordances. In: Shaw R, Bransford J, eds, Perceiving, Acting, and Knowing: Toward an Ecological Psychology. Hillsdale, NJ. Lawrence Erlbaum. 1977. 67–82.
[19] Glasgow J, Narayanan NH, Chandrasekaran B, eds. Diagrammatic Reasoning: Cognitive and Computational Perspectives. Menlo Park. AAAI Press, 1995.
[20] Hobbs J. Granularity. IJCAI. 1985. 432–435.
[21] Lakafosis V, Yi X, Le T, Gebara E, Wang Y, Tentzeris MM. Wireless sensing with smart skins. Sensors. IEEE. 2011. 623–626.
[22] Lakoff G, Johnson M. Metaphors we live by. University of Chicago Press. Chicago. 1980.
[23] Lungarella M, Hafner V, Pfeifer R, Yokoi H. An artificial whisker sensor for robotics. IEEE Conf on Intelligent robots and systems. 2002. 2931–2936.
[24] Marr D. Vision. Cambridge, MA: MIT Press, 1982.
[25] Nebel B, Bürckert HJ. Reasoning about temporal relations: a maximal tractable subclass of Allen's interval algebra. JACM, 1995, 42(1): 43–66
[26] Mossakowski T, Moratz R. Qualitative reasoning about relative direction of oriented points. Artificial Intelligence. 2012. 180. 34–45.
[27] Norman DA. The Psychology of Everyday Things, New York: Basic Books, Inc. 1980.
[28] Palmer SE. Fundamental aspects of cognitive representation. In: Rosch E, Lloyd BB, eds. Cognition and categorization. Hillsdale. Lawrence Erlbaum. 1978. 259–303.
[29] Pfeifer R, Scheier C. Understanding Intelligence., Cambridge, MA: MIT Press, 2001.
[30] Schultheis H, Barkowsky T. Casimir: An architecture for mental spatial knowledge processing. Topics in Cognitive Science. 2011, 3(4): 778–795.
[31] Schultheis H, Bertel S, Barkowsky T. Modeling mental spatial reasoning about cardinal

directions. Cognitive Science, 2014, 38(8): 1521–1561.

[32] Sloman A. Why we need many knowledge representation formalisms. In: Bramer M, ed. Research and development in expert systems, New York: Cambridge University Press, 1985: 163–183.

[33] Spenneberg D, Kirchner F, 2007. The Bio-Inspired SCORPION Robot: Design, Control & Lessons Learned. In: Zhang H, ed. Climbing & Walking Robots, Towards New Applications, Itech Education and Publishing, Vienna. p 546 ff.

[34] Tanji J, Takagi T, Qiu J. Intelligent material systems: Application of functional materials. Appl. Mech. Rev., 1998, 51(8): 505–521.

[35] Tenorth M, Beetz M. KnowRob – A knowledge processing infrastructure for cognition-enabled robots. International Journal of Robotics Research (IJRR), 2013, 32(5): 566–590.

[36] Wintermute S, Laird JE. Bimodal spatial reasoning with continuous motion. Proc AAAI. 2008. 1331–1337.

[37] Freksa C. Strong spatial cognition. Conference on Spatial Information Theory XII(COSIT 2015). Lecture Notes of Computer Science. Springer 2015.