Social-ecological modelling with LARA: A psychologically well-founded lightweight agent architecture

Ramón Briegel, Andreas Ernst, Sascha Holzhauer, Daniel Klemm, Friedrich Krebs, Aldo Martínez Piñánez
University of Kassel, Center for Environmental Systems Research (CESR)
{briegel, ernst, holzhauer, klemm, krebs, martinezpinanez}@usf.uni-kassel.de

Abstract: For the purpose of policy simulation in coupled social-ecological systems (e.g. energy supply), a credible modelling of actors – especially citizens – and their decision processes is needed. This requires a framework capable of handling high numbers of heterogeneous agents (several hundreds of thousands). In our presentation we describe a framework called LARA (Lightweight Architecture for boundedly Rational Agents) which meets these requirements and fills the gap between frameworks without built-in psychological foundations and full-fledged cognitive architectures which are both not viable options in this context. LARA provides prefabricated components of an agent's decision process like perception, memory, and different modes of decision making. These components are psychologically plausible, i.e. based on appropriate psychological results and theories. Moreover, interfaces for basic learning and social influence are available. To demonstrate the flexibility of LARA, we present an exemplary application model.

Keywords: Agent based modelling; social-ecological modelling; decision modelling; agent architecture; bounded rationality.

1   INTRODUCTION

In the modelling of coupled social-ecological systems like in various other domains, agent-based modelling has proven to be a useful method for investigating the behaviour of complex systems consisting of heterogeneous agents, i.e. autonomous entities which interact with each other and with their environment (Grimm et al [2006], Gilbert & Troitzsch [2005], Railsback [2011]). Despite the fact that social-ecological systems include human agents, existing agent-based models (ABMs) of such systems often lack psychological foundations in their way of mapping human behaviour (Ernst [2009]). In contrast, other models of human behaviour *do* take account of empirical and theoretical psychological insights; they rely on highly complex cognitive agent architectures designed for a specific application context. Such models, however, are difficult to generalize or to be transferred to other contexts. Furthermore, implementing such a sophisticated model is a cumbersome task that requires a lot of time and technical programming capabilities.

To escape this dilemma, a framework is needed which on one hand is based on psychological findings and theories and on the other hand is usable in a wider range of modelling contexts and simplifies the implementation of ABMs of social-ecological systems. Especially in the field of policy simulation where a credible modelling of citizens is a crucial element for estimating the effects of political measures in the population, such a framework could enhance the usability and the reliability of ABMs. The LARA framework (Lightweight Architecture for boundedly Rational Agents) has been designed to meet these requirements and to fill the gap between frameworks without built-in psychological foundations and specialized full-fledged cognitive architectures.

Our paper is organized as follows: The LARA framework, its components and processes together with an example model are described in section 2. Section 3 explains how to use LARA, i.e. the steps which are necessary for implementing a concrete model based on LARA. The results of the example model are presented in section 4. We conclude in section 5 by pointing out the value of LARA for future implementations of social-ecological models.

2   DESCRIPTION OF THE LARA FRAMEWORK

2.1   Overview

Built in a modular structure, the LARA framework consists of largely independent prefabricated components which interact to map human agents' information processing, decision making and behaviour. The agents can interact with one or more environments, e.g. biophysical, socioeconomic, and social environment. These environments may be subject to exogenous and/or endogenous dynamics. The time concept is based on discrete time steps.

Figure 1 shows an overview of the architecture and an agent's standard path of processing in each simulation time step. The environments are perceived by the agent (according to its subjective information filtering); these perceptions are stored to the agent's memory. The agent's fundamental goals[1], behavioural options and preference structure (goal weights) are also represented in the memory. The pre-processor then selects an appropriate decision mode (habit, heuristics, deliberative evaluation) and the set of currently feasible behavioural options. The actual selection of a behavioural option is performed by the action selection component. Finally, the post-processor component is responsible for evaluation, storing and possible adaptations of the selected behavioural options.
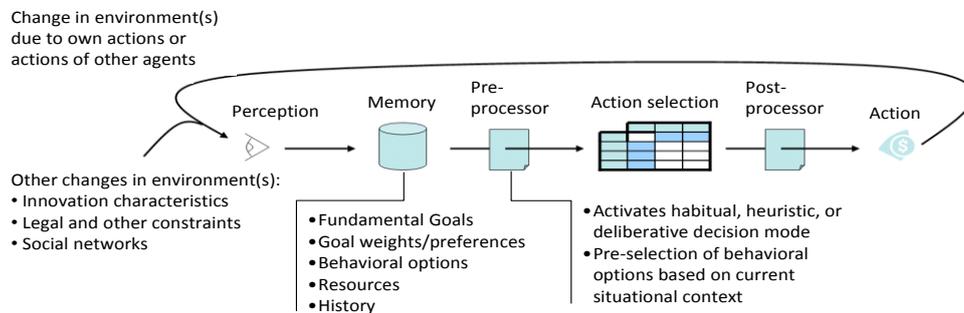
Figure 1: Main components of LARA and their function

2.2   Example application: The houseplant model

We explain LARA's basic features by modelling the growth of two indoor plants located at two common rooms in two different flats. Three agents (Leon, Mia, and Lukas) share a flat in Wilhelmshöher Allee 1 and two other agents (Ben and Hanna) in Herkulesstraße 2. As its social environment, each group has a unique collaborators network consisting of the other household members. The agents dispose of two behavioural options: "Irrigate" or "Do Nothing"; their goals are "Having a Nice Plant (HNP)" and "Being Social (BS)." Figure 2 shows a schematic representation of the model and the goals' weights (cf. section 2.6) for each agent.

---

[1] By "fundamental goal", we mean anything more or less abstract which the agent aims at in principle or in the long term. This subsumes notions and concepts like need satisfaction, motive, value, orientor, and ideal. For sake of simplicity, we will use "goal" as a synonym of "fundamental goal" in this paper.
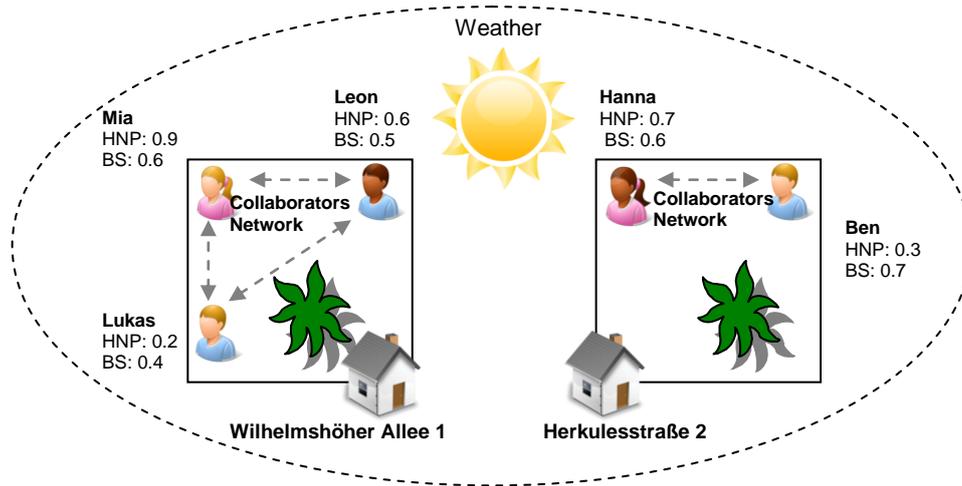
Figure 2: Schematic Representation of the Houseplant Model

In every time step (which represents one day), the environment is updated by setting the sunshine duration. If an agent has irrigated the plant, a fixed amount is added to the water stock of the plant. As well, depending on the sunshine and the plant's size, water stock is reduced as growing plants consume more water. Subsequently, the size of every plant is updated. If its water stock is enough, the plant grows with a different velocity according to its size: a small plant grows faster than a bigger one.

2.3 Perception

Perception is the basic process of translating physical information external to the individual into a format that lends itself to further internal processing. Such an internal representation (of a particular attribute of an environment at a particular point in time) is generally called a percept. Environments can be classified into the three categories of biophysical, i.e. natural or built-up environments, socioeconomic environments, i.e. political, legal or economic constraints, and social environments, i.e. the network that enables social exchange and appraisal between agents (Ernst et al [2009]).

A main characteristic of human perception is its strong selection and aggregation component. In everyday contexts, the wealth of information is significantly reduced during the translation process. This reduction is steered by the attention given to a specific item, and depends both on top-down individual motivation and characteristics of the physical signal. The more realistic an agent-based model becomes and the more data it encompasses, the more important such attentional processes become.

Qualitative and quantitative perceptions can be transformed into symbolic, internal representations, e.g. by some form of evaluation or appraisal (threshold values), which are then aggregated appropriately and stored in memory.

In our example model, every agent perceives the sunshine of the current day, the water stock of the plant, and the percentage of collaborators that did irrigate in the last time step.

2.4 Memory

Elements stored in human memory either stem from the perception of external signals or are the result of cognitive, internal operations. Many of the behavioural phenomena discussed as bounded rationality are related to processes of human memory that make it differ from a simple, yet perfect data repository. E.g., forgetting of elements stored in long-term memory is dependent on the character of

the elements, with highly connected and significant elements being retained longer than isolated or insignificant elements.

Together with the observed element, humans usually retain (part of) the context of that perception together with it. E.g., not only a story is remembered, but where it was told as well. This stored situational context also proves crucial in recalling elements from memory when needed.

As a general architecture, LARA provides a standard implementation to represent these basic phenomena.

The entries of the agents' memory – which may be quantitative and/or qualitative – are formed on the basis of current percepts each of which being composed of the following elements:

- the percept itself,
- the current context,
- the time the entry was memorised, and
- possibly the retention time of the entry.

In general, a context represents a relevant subset of the state of the agent at the time the perception occurred. Therefore, entries can be marked with a key that represents the according context adequately and possibly hierarchically and enables the retrieval of all entries for a particular context.

The retention time of an entry – which may be different for different percepts and contexts – is set when the entry is created. It determines how long the entry will remain in memory. Moreover, the architecture allows for entries which are not directly grounded in perception such as cognitions about causal relations or inferences etc. as may be generated by the post-processor.

In addition to storing new entries, the memory component is used in different ways throughout the various LARA components like e.g. collecting experience, forming habits or learning. Details are documented in the following sections.

2.5   Decision modes, decision trees and the preprocessor

Humans are able to tune their behavioural responses to the current situational necessities and the resources they have. We distinguish three types of responses which we call the habitual, heuristic, and deliberative decision mode, respectively.

Most of daily behaviour seems to be habit driven (Aarts & Dijksterhuis [2000]), e.g. mobility behaviour or water and energy use. This represents a "fast lane" of behaviour that is useful when positive previous experiences with a situation exist and a known behavioural option can be readily applied.

If this is not the case, current situation and available behavioural responses have to be elaborated and matched in a conscious, reflexive process. This process has to rely on knowledge about the utility of the behavioural options in the current situation – knowledge which has to be gathered beforehand or generated by mental simulation processes. In decision theory, this is often called goal-directed, multi-attribute utility decision making (Baron [2000]).

Where such deep knowledge cannot be provided, humans tend to apply rules-of-thumb, so-called heuristics (Gigerenzer et al [2001]). They operate on a very restricted set of available pieces of knowledge, but provide however suitable behavioural responses for a wide range of situations.

The preprocessor of LARA has the tasks of determining the decision mode for a given situation and of setting the context for the process of action selection which follows. As depicted in Figure 3 the pre-processor encompasses up to five steps.

The decision mode selector assigns the way the agent selects the behavioural option(s) to be performed with respect to the demands of a particular situation. LARA currently supports deliberative decision making, decision trees, and habit. For instance, the mode selector can evaluate previously executed behavioural options (BOs) and use a BO as a habit for a certain context if it was executed repeatedly and successfully before. Some or all of the following pre-processing steps then become obsolete and do not need to be executed.

The BO collector retrieves behavioural options from the agent's memory. Here, the number of considered BOs may be limited by retrieving only those that were updated during the last x time steps. The BO pre-selector filters out context-

irrelevant behavioural options. The default pre-selector retains only those BOs that do have at least one non-negative utility for any of the relevant fundamental goals.

Updating of a behavioural option's utility values is triggered by the BO updater. It is also responsible for quantifying those entries in memory that are represented qualitatively in order to enable the evaluation of those behavioural options. As default, the BO updater delegates this task to each of the BOs. Finally, the goal weights are updated by the preference updater.

If needed for a special model purpose, the above described order of execution may be modified by the decision mode selector.
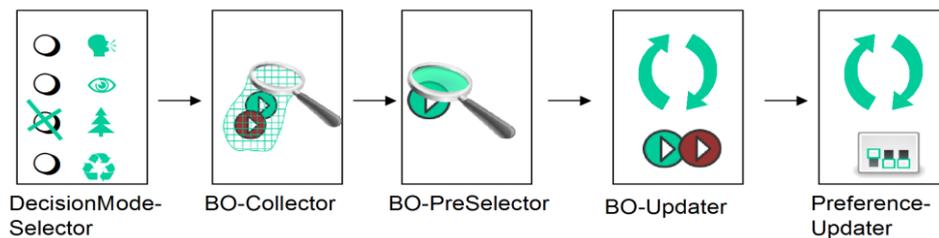
Figure 3: These steps are usually performed in the pre-processing stage

Decision trees are provided by LARA in order to model rule-based decision making. The modeller builds up the tree by choosing left and right branches and defining an evaluate method that decides which sub-tree to follow. The leaves of the decision tree represent the selected BO(s).

The habit decider automatically selects the behavioural option which was chosen the last time, whereas the exploration decider selects a BO from the pre-processed set at random.

In our example model, if the plant's water stock is appropriate, agents act according to their habits and deliberation does not take any role. However, if water stock of the plant is not appropriate, every agent deliberates about irrigating a plant or not.

2.6 Action selection and situational utility

After the stage has been set by pre-processing the situational information from perception, memory, and internal states, an appropriate decision has to be made. While habitual behaviour consists of a mere "same as before" and heuristics are represented as decision trees in the pre-processor, deliberative decision making combines characteristics of the current situation and individual aspects of the decision maker.

We do suppose that people deliberately pursue their goals, whatever may be the nature of these goals, if there are degrees of freedom to do so. The core elements for deliberate decision making encompass

(a) a decision maker's goal preferences,

(b) a set of behavioural options from which a response appropriate to a given situation can be chosen, and

(c) some knowledge about the utility of some behavioural option to one or more of the decision maker's goals, thus linking both (a) and (b) (Ernst [2003]).

This knowledge is acquired over the course of action and serves to build expectations about outcomes of own behaviour. The specific ways of weighting and aggregating such situational knowledge may vary from situation to situation and from person to person.

Crucial to the realism of a citizen agent is its capability to adapt its behaviour to changing circumstances. While a gradual low level behavioural adaptation can be readily represented in the LARA framework, higher level learning requires artificial intelligence mechanisms that are not an integral part of the architecture presented here. In any case, human learning stems from the evaluation of the satisfaction with the outcomes of one's actions in a given situation. Indicators of goal attainment,

utility of behavioural options, and of competence to master certain situations should be available to be stored in memory and serve as a basis for further actions (Ernst [2003]).

The agent's beliefs concerning the effectiveness of its available behavioural options for the pursuit of its fundamental goals are represented in LARA as the agent's basic utility matrix. For each pair of behavioural option and goal, the corresponding matrix element measures how appropriate this behavioural option seems to the agent for attaining this goal. Negative values are allowed (indicating that this behavioural option is detrimental for the attainment of that goal).

Each of the agent's various goals is associated a number – the goal's weight – which represents the basic subjective importance of that goal for the agent.

Due to situational circumstances, the current, effective goal weights of an agent may differ from the above mentioned basic weights. For example, the current value of a state variable (of the agent or of one of its environments) may boost the urgency to attain a certain goal, or a specific situation may offer an opportunity for attaining a certain goal, which therefore becomes temporarily more important. This phenomenon may even reverse the agent's goal preferences in a particular time step. The situational impact on goal weights is mapped by goal-specific factors; to obtain the situational goal weight, the basic goal weight is then multiplied by the corresponding factor.

If the agent decides in deliberative mode, the agent's actual selection of a behavioural option will then be based on the situational utility matrix, which is obtained by multiplying each element of the basic utility matrix with the situational weight of the corresponding goal. So this matrix reflects the partial utilities of the various behavioural options in light of the goal importances under the current conditions (represented by the situational goal weights). Finally, a deliberative choice component (chosen by the modeller) selects one or more BOs for execution; e.g., the maximum sum deliberative choice component selects the behavioural option with highest sum of situational partial utilities.

In the houseplant model, under the deliberative decision mode, the behavioural option "Irrigate" is more valuable for the agent when the agent has a high preference for having a nice plant and being social and/or if the plant has a strong need for water. On the contrary, too much water makes this behavioural option less attractive. The utility of the behavioural option "Do Nothing" is set opposite to the utility of "Irrigate" in an obvious way.

2.7   Post-processing

Sometimes the mere fact of having taken a decision for or against a particular course of action can affect ones appraisal of that option or that decision. It can be observed that decisions, once made, are immunized against conflicting evidence and kept a certain time (Montgomery [1987]).

LARA provides an interface for such (possibly irrational) post-processing of decisions. The post-processor component is triggered after deciding and before performing the action in order to evaluate the decision and store selected behavioural options in the agent's memory. It can be easily extended by user-defined classes that incorporate behaviour as mentioned above. Furthermore, the architecture's flexibility allows the modeller to integrate further steps, e.g. updating the BO's utility after taking action in the sense of simple learning.

3   FRAMEWORK IMPLEMENTATION

The framework's design reflects the pursuit of flexibility, light weight and performance. Its architecture is strongly component-based. Each aforementioned component (perception, memory, pre-processor, decision making, post-processing) is independent, and coupling is realised by an event-bus: Subsequent components subscribe at the event-bus for certain events. The current component then publishes a certain event, and the event-bus triggers subscribed instances for that particular event. Thus, the event-bus supports flexible flow control and state-dependent activation as well as ordering of component execution. It allows the

definition of depending events that need to be executed before another one. Furthermore, parallel execution of tasks becomes easier. As each decision is identified by a decision configuration object, it is possible to define several decisions for potentially different agent groups that are executed sequentially or in parallel.

The framework is programmed in Java and thus ensures platform independence. As an agent framework it may be used stand-alone or integrated in every Java-based ABM-framework, for instance Repast Simphony (North [2007]). Since LARA allows the attachment of required code on the agent side as an agent component it can be easily linked to existing agent-based models.

Many parts of the components use fine grained interfaces with default classes that may be exchanged by user-defined realisations of that interface (e.g. the decision mode selection class). The incorporation of Log4J as an accurately configurable logging framework and the use of a random number manager facilitate the analysis of agent-based models developed with LARA.

4   RESULTS OF AN EXAMPLARY SIMULATION RUN

The results of a simulation run for the example model are summarized in Figure 4 below. The upper first section of the figure shows sunshine duration within one solar year on a daily basis taken from sample observations. In the second part of the figure, it is plotted how the agents have chosen to act in Wilhelmshöher Allee 1 and Herkulesstraße 2 according to their goals and preferences; that is, when they decide either to irrigate or not. Below the representation of agent behaviour, it is shown how the water stock of each plant fluctuates as it is affected by sunshine duration, behaviour of the agents, and actual plant's size. Finally, the development of the plants size during the simulation is shown.
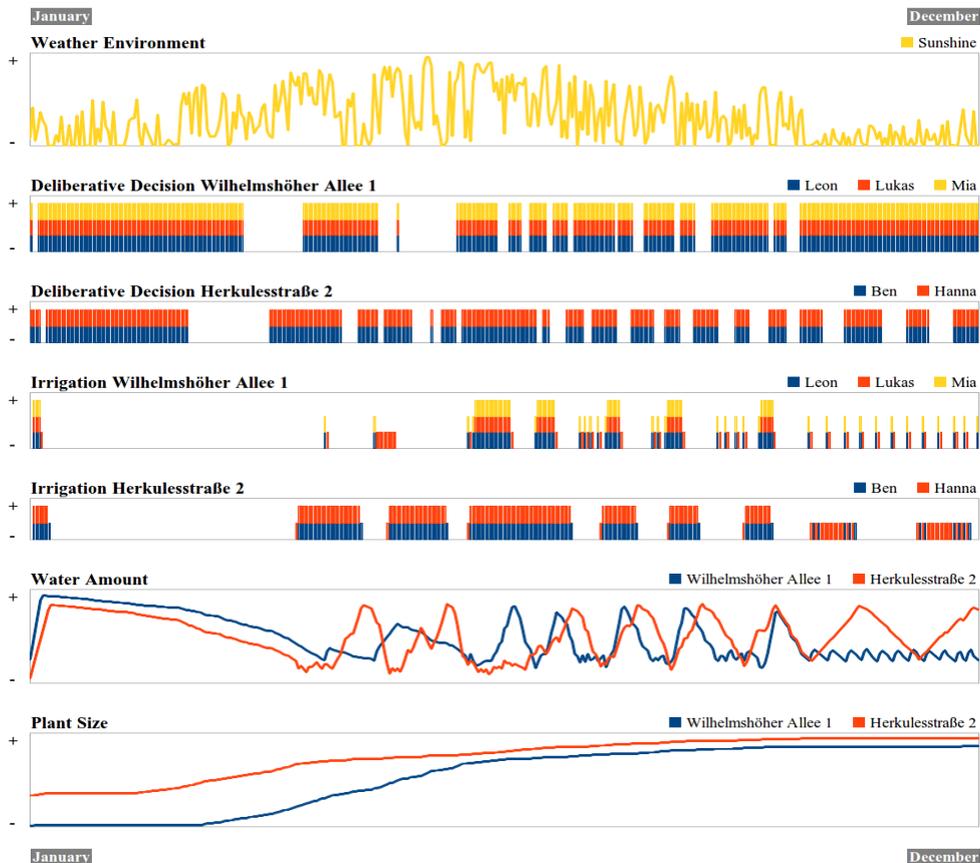
Figure 4: Summary of Simulation Results

In the very beginning the agents in both communities irrigate the plant, which then becomes habitual. When the water amount reaches a too high level, the agents make a deliberative decision every day until the water level drops below the critical level. Then to not water the plant becomes habitual until the water amount gets too low as the sunnier period of the year begins. Subsequently this cycle repeats in intervals of changing duration, depending on the amount of sunshine and the diversity of the communities.

5   CONCLUSION

We have presented a psychologically well-founded framework for agent-based modelling of social-ecological systems which is capable of handling large numbers of heterogeneous agents and provides different situation-specific modes of decision making. The LARA framework fills the gap between ABM frameworks without psychological foundations and highly specialized cognitive architectures which have been designed for a very specific application context.

The architecture's component-based nature in combination with the event-bus approach allows the modeller to easily interchange pre-defined LARA components with classes adapted to the application model's needs. Thus, LARA is usable in a wide range of modelling contexts, especially in policy modelling, and facilitates an appropriate implementation of decision models. At CESR, LARA is currently being used in two projects: SPREAD (Scenarios of Perception and Reaction to Adaptation), which investigates the spreading of socio-technological innovations in the renewable energy sector like switching to green electricity and investing in community managed solar plants), and KUBUS (Supporting regional climate change adaptation by means of socio-environmental surveys and scenario development), where simulations are used to better understand the public reaction to policy-defined adaptation strategies via the individual processes of adaptation to climate change.

REFERENCES

Aarts, H. & Dijksterhuis, A., Habits as knowledge structures. Automaticity in goal-directed behavior, *Journal of Personality and Social Psychology*, 78(1), 53–63, 2000.

Baron, J., *Thinking and Deciding*, Cambridge University Press, Cambridge, MA, 2000.

Ernst, A., Agentenbasierte Modellierung des Handelns in Gemeingutdilemmata, *Jahrbuch Ökologische Ökonomik*, 3, 139–170, 2003.

Ernst, A., Interaktion, Dynamik, Raum – Komplexe agentenbasierte Modelle in der Umweltpsychologie. *Umweltpsychologie*, 13(1), 84–98, 2009.

Ernst, A., Krebs, F., Elbers, M., Holzhauer, S. & Klemm, D., LARA: A lightweight architecture for boundedly rational citizen agents, long abstract submitted to the 6th conference of the European Social Simulation Association (ESSA 2009), 2009.

Gigerenzer, G., Todd, P.M. & the ABC Research Group, *Simple heuristics that make us smart*, Oxford University Press, Oxford, 2001.

Gilbert, N. & Troitzsch, K.G., *Simulation for the Social Scientist*, Open University Press, 2005.

Montgomery, H., Image theory and dominance search theory: how is decision making actually done?, *Acta Psychologica*, 66, 221–224, 1987.

North, M., Howe, T., Collier, N., Vos, J., Takahashi, S., Sallach, D. & Rouchier, J. (Eds.), Declarative Model Assembly Infrastructure for Verification and Validation Advancing Social Simulation: The First World Congress, Springer, Heidelberg, 2007.

Railsback, S.F. & Grimm, V., Agent-based and individual-based modeling: A practical introduction, Princeton University Press, Princeton, NJ, 2011.