

From *Inter-Agent* to *Intra-Agent* Representations: Mapping Social Scenarios to Agent-Role Descriptions

Giovanni Sileno¹, Alexander Boer¹ and Tom Van Engers¹

¹Leibniz Center for Law, University of Amsterdam, Netherlands

{g.sileno, a.w.f.boer, t.m.vanengers}@uva.nl

Keywords: Agent-roles, Scenario-based modeling, Social systems, Institutions, Story animation, Narratives

Abstract: The paper introduces elements of a methodology for the acquisition of descriptions of social scenarios (e.g. cases) and for their synthesis to agent-based models. It proceeds along three steps. First, the case is analyzed at signal layer, i.e. the messages exchanged between actors. Second, the signal layer is enriched with implicit actions, intentions, and conditions necessary for the story to occur. This elicitation is based on elements provided with the story, common-sense, expert knowledge and direct interaction with the narrator. Third, the resulting scenario representation is synthesized as agent programs. These scripts correspond to descriptions of agent-roles observed in that social setting.

1 INTRODUCTION

Research in AI and computer science usually studies normative systems from an engineering perspective: the rules of the system are designed assuming that the behaviour of the system and of its components is mostly *norm-driven*. In human societies, however, the position of policy makers and regulators is completely different. The target social system exists and behaves around them, and they are part of it. Moreover, human behaviour is actually *norm-guided*. People adapt to their social environment, both influencing and being influenced by institutions. Multiplicity of institutional conceptualizations and non-compliance (intentional or not) are thus *systemic*.

We presented in previous works (Boer and van Engers, 2011b; Boer and van Engers, 2011a; Sileno et al., 2012) elements of an application framework based on legal narratives, such as court proceedings, or scenarios provided by legal experts to make a point about the implementation or application of the law.¹ In this, we are investigating a methodology for the acquisition of computational models of such interpretations of social behaviour. How to represent what people know about the social system, or, equivalently, how people (re)act in a social system? Our focus is

¹These narratives are particularly interesting because they are *produced* by the legal system, with the intent of transmitting - within its current and future components - relevant social behaviours and associated institutional interpretations.

not on the narrative object, but on the knowledge that observers and narrators handle, when they observe social behaviour and generate explanations. The assumption of *systemicity* is thus not related to the discourse, but to a cognitive level.

Reducing the problem to the core, this work investigates the transformation of a sequence of *inter-agent* interactions in *intra-agent* characterizations, reproducible in a computational framework. In section 2, we present our case study. We analyze it at a *signal layer*, defining the topology and the flow of the story. In section 3 we show how to enrich the previous representations with an *intentional layer*, integrating *institutional concepts* as well. In section 4 we provide elements about the transformation of the previous models into scripts for cognitive agents. Discussion ends the paper.

2 INTER-AGENT DESCRIPTION

Despite its simplicity, a short story about a sale transaction provides a good case of study.

A seller offers a good for a certain amount of money. A buyer accepts his offer. The buyer pays the sum. The seller delivers the good.

A successful sale is a fundamental economic transaction. Consequently, what the case describes is a behavioural pattern used both in the performance and in the interpretation of many other scenarios.




Figure 1: Message sequence chart of a story about a sale.

2.1 Signal layer

The story describes four events, namely four acts performed by two agents.² The first two acts are easily recognized as speech acts, but all these actions can be *teleologically* interpreted as bringing some informational change into another agent. From this perspective, we can consider all of them as acts of communications, i.e. as messages going from a sender to a receiver entity. Thus, the previous story can be illustrated using a communication diagram, for instance as a *message sequence chart* (MSC).³ Simplifying the notation, we obtain the illustration in Fig. 1.

So far all seems easy, but the sale process described before lacks some important details. For instance, in a marketplace, paying and delivering are physical actions. They produce some consequence in the world: money, goods move from one place to another. Furthermore, payment and delivery are acknowledged by perception. Second, the buyer has accepted the offer, because he was somehow receptive toward that kind of messages. Third, a buyer who already paid usually does not leave without taking the good, just as the seller does not allow the buyer to go away with an object without paying. In our story, all goes well, so the narrative does not provide any element concerning these checks. However, this does not mean necessarily that the buyer and seller had not checked if everything was fine. Fourth, sometimes a buyer takes the good and then pays, sometimes the order of actions is inverted. These four points reflect characteristics which are left implicit in the story, and, consequently, in the MSC:

- acts have side-effects on the environment (at the very least, a transient in the medium transferring the signal),
- an action consists of an emission (associated to an

²There is also another implicit agent, the narrator, but he will be neglected in this work, supposing he has a pure descriptive intention.

³MSCs are the basis of the *sequence diagrams*, one of the behavioural diagrams used in UML. For further information, see for instance (Harel and Thiagarajan, 2004).




Figure 2: Topology of the story.

agent) and of a reception (associated to a *patient*),

- certain actions have a *closed-loop* control: agents perform some monitoring on expected outcomes,
- the sequence of events/acts in a story is often a partial order, hidden by the linear order of the discourse.

Apart from the last point, which could be solved with the UML “par” grouping for parallel constructs, we have to find alternative representations to help the modeler in scoping and refining the content of the story.

2.1.1 Topology of the story

Inspired by the Actor model (Hewitt et al., 1973), we have drawn in Fig. 2 the *topology* of the story. The topology serves as a still picture of the whole case, and show how signals are distributed between the characters. The little boxes are *messages queues*, the lines identify *communication channels*. The story describes which specific *propositional content* is used in the exchanged messages. In order to take eventual side-effects into account, we introduced an explicit “world” actor, disjoining the emission from the reception. The optional part of the communication is visualized with dotted lines. The world would play as intermediary entity also in case of broadcast messages.⁴

2.1.2 Flow of the story

Orthogonal to the topology, we define the flow of the story as the order in which events occurred. As a first definition, we may consider a story as a chain of events (a strictly ordered set):

$$\mathcal{E} = \{e_1, e_2, \dots, e_n\} \quad (1)$$

In narratology this layer is usually called the *fabula*: “a series of logically and chronologically related events [..]” (Bal, 1997). This name dates back to Propp, which, altogether with the Russian formalists,

⁴With a similar spirit, communication acts performed autonomously by the world actor can model natural events.

started considering each event in the story as *functional*, i.e. a part of a whole sequence, necessary to bring the narrated world from initial conditions to a certain conclusion. Furthermore, specific circumstances may be described in correspondence to the occurrence of an event. As a result, a story corresponds to the following chain:

$$C_0 \xrightarrow{e_1} C_1 \xrightarrow{e_2} \dots \xrightarrow{e_n} C_n \quad (2)$$

where e_i are associated to transitions and C_i is a set of conditions assumed to continue at least until the occurrence of e_i .

Consequence and consecutiveness This definition may look very simple, but the manifold relations between *consequence* (logical, causal, ..) and *consecutiveness* (informed by time, ordering, ..) are actually very delicate to assess. Furthermore, two different chronological coordinates coexist in a narrative: a *story-relative* time, i.e. when the event has occurred in the story, and a *discourse-relative* time, i.e. when that event has been reported or observed.

In order to unravel this knot, we use a four steps methodology to reconstruct the relations between the elements of the story.

First, we elicit relevant abstractions which are used in the interpretation. In particular, we define an event/condition as *free* if the interpreter does not acknowledge any relation⁵ with another event or condition in the story. We refer to such relations as *dependencies*. Some dependencies are *syntactic*. For instance, you can accept an offer only if there is an offer, i.e. if an offer has been previously made. Others are *contextual to the domain*. For instance, in a web sale, payment usually occurs before delivery. In all cases, dependencies can be used to put a *strong constraint* on the ordering of events.

Second, there may be clues of the story-relative time within the text. Time positions and durations are usually meant to give some landmark to the listener. They are described in absolute or relative terms. When a listener interprets them, it creates a relation between events, *contingent to the story*. Such relative positioning constitute the *medium constraint*.

Third, if we have no clues about dependencies, or temporal relations between events, a possible sequence is at least suggested by the discourse-relative time. This provides a *weak constraint* on the ordering of free events.⁶

⁵Apart having occurred in the same story.

⁶The *story* and *discourse contingencies* of the medium and weak constraints become *contextualities* if they are entailed by strong constraints.

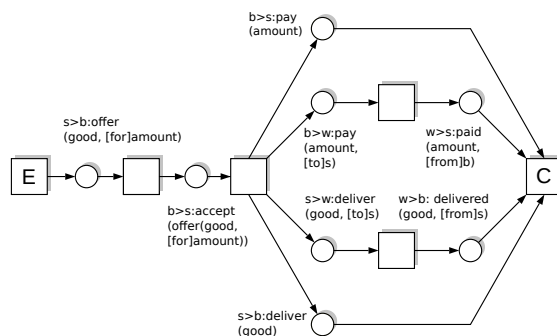


Figure 3: Flow of the story.

If all three constraints are satisfied, we do not expect any concurrent events, at least within one story frame.⁷ However, it is easy to object to such a strict determination.

Consequently, at the fourth step, we weaken the previous strict temporal constraints (e.g. from $e_{i+1} > e_i$ to $e_{i+1} \geq e_i$) in two cases: (a) dependencies can be associated to no-time-consuming processes (e.g. logical equivalences); (b) events may occur simultaneously, when triggered by parallel sub-systems. Furthermore, the medium and weak constraints refers to *contingent* relations (according to the modeler). In order to be able to compare the internal structure of stories, we can neglect them. With these modifications, the set \mathcal{E} defined in (1) is a partially ordered set.

Let us take our case story. There is a relation of syntactic necessity between offer, acceptance and performance. In addition, there are two agents. These entities can be considered as parallel systems, that may concurrently interact with the world. Therefore, without further contextual specification, payment and delivery are concurrent events.

Visualization A simple way to visualize the flow of the story is by the use Petri nets, as we did in Fig. 3. We opted for a practical naming of places: *sender>receiver:content*. At this point, places represent messages, associated to speech acts. Actions are like “compacted” into transitions.

The main scope of the flow is to preserve the *story synchronization*. Further layers may be integrated, increasing the granularity of the description, maintaining the previous points of synchronization, in the same spirit of *hierarchical Petri nets* (Fehling, 1993). For instance, in Fig. 3, we have disjoint the generation of the message from its reception using an intermediate “world” actor, as we did in the topology.

⁷A more complex story may consist of many frames. For “frame”, we consider a sub-story that follows the Aristotelian canon of unities of time, space and action.

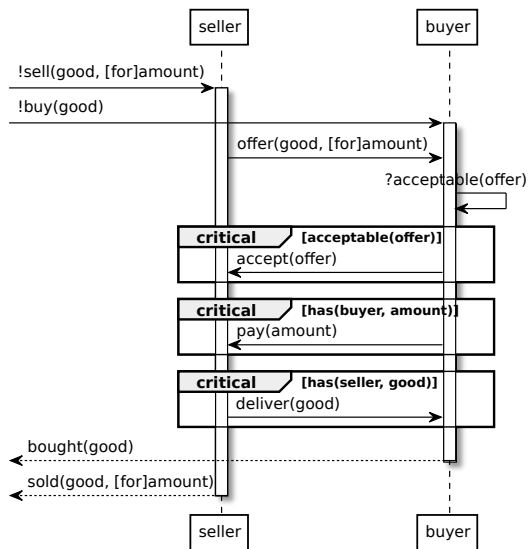


Figure 4: MSC with intentions and critical conditions.

3 AGENTIC CHARACTERIZATION

In the previous section we referred to the messages exchanged in a social scenario, and narrated through a story. When we interpret such messages, however, we apply an intentional stance just as we do in our experiential life: we read the actors as intentional agents, attributing them beliefs, desires and intentions. This *ex-post* intentional interpretation of the story results in a decomposition of the plans followed by the agents. In addition, in order to trigger or enable the performance of the reported acts, there may be other relevant conditions or hidden acts to be taken into account. They could have been left implicit by the narrator, but plausibly they are known to the target audience of that narrative. This is the main assumption on which our work is based: any reader/modeler can always provide one or several reconstructions of what happens behind the signal layer, supported by common knowledge and domain experts (or, at least, interviewing the narrator himself, if needed).

3.1 Acquisition methodology

As we did before, we start the acquisition using a MSC diagram. A possible outcome of the interpretation is in Fig. 4, where we introduce adequate extensions.

First, we consider externalized intents (with a “!” prefix) as the events triggering the processes of buying/selling. The final outcomes of those actions are then reported with output messages as well, at the end

of the chart. Second, we add eventual hidden acts. In our case, we know that a buyer usually accepts an offer only after positively evaluating it.⁸ Third, we use the *critical* grouping to highlight which conditions (in addition to sequential constraints) are necessary for the *production* of that message.

To sum up, in our story, we add that: (a) the buyer performs an evaluation of the offer (*evaluation* action), (b) the buyer accepts the offer if it is acceptable for him (*acceptability* condition), (c) the buyer pays (the seller delivers) if he owns the requested money (good) (*ownership* condition).

The MSC diagram in Fig. 4 furnishes a good summary of the story: the inputs/outputs provide an intentional characterization, the vertical bars indicate the ongoing activities, while the messages refer to successful acts of emission and reception, whose occurrence is constrained by the critical conditions. Unfortunately, further refinement is necessary to cover the basic figures encountered in an operational setting. There is no separation between emission and reception, and between epistemic and ontological. For instance, we cannot distinguish between a case in which the buyer thinks he has not enough money, and another in which he thinks that there is enough on his bank account, but the bank does not “agree” with him.

A simple solution to this problem would be to add intermediate actors (e.g. the bank), localizing where the failure occurred. In a complex case, however, the resulting visualization may be overloaded. In the following sections we will therefore introduce some patterns to be attached to the flow of the story. Instead of using just one visualization, our approach aims to provide alternative representational cuts.

3.1.1 Hierarchical tasks

First, we elicit the hierarchical decompositions of activities performed by the agents. These serve as basic schemes for the behavioural characterization of the agents, and use hierarchical, serial/parallel constructs. In practice, this is obtained by identifying, in the story flow, the activities of the actor as *agent* (emitter) and *patient* (receiver), and relating them according to their dependencies. Fig. 5 reports the result of this step for the buyer.

3.1.2 Emission and reception

Second, activities are anchored to messages. This is a delicate phase: we want to maintain the synchro-

⁸The absence of evaluation is symptomatic of a *combine* scheme: the buyer performs mechanically the acceptance in order to advance the interests of another element of his social network.




Figure 5: Hierarchical decomposition of tasks.

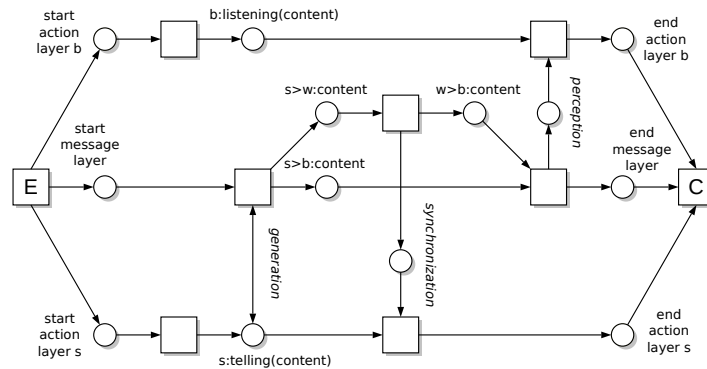


Figure 6: Communication pattern.

nization given by the story *and* the dependencies associated to the activities. Fig. 6 reports our solution (applied on a single message) which explicitly divide emission from reception. The proposed Petri net is complete and well-formed. It is scalable to multiple agents, adding a reception cluster (e.g. $w>b:content$, *perception*, etc.), in order to connect the message to each agent that is reachable by the communication.

Note: partial orderings may hold independently in the story flow *and* in the activity diagrams. In this case, for instance, we do not know a priori if payment occurs before delivery (as acts), as we do not know if the buyer pays before monitoring the delivery (as actions), or vice-versa, or simultaneously.⁹

3.1.3 Illocutionary acts

Third, we recognize the practical effects of messages. Beside of being signal (or a *locutionary act*), each message is associated to an *illocutionary act*, and then, when it is put in a computational form, it should integrate some *pragmatic* meaning.

For simplicity, we consider only four types of performatives: *assertions*, *commissives*, *commands* and

⁹The simultaneity of actions is evident when considering a collective agency. However, considering an individual physical agent, we can assume at least that, concurrently to the performance of a definite activity, he maintains some listening activities as well.

inquiries. Moreover, we interpret commitments as obligations to the self and commands as attempts to instill obligations into the receiver. Considering our story, we have:

- the offer is a conditional promise: the seller *commits* to deliver the good to a buyer who commits to pay his price;
- the acceptance is a promise: the buyer *commits* to pay a given amount;
- the payment can be interpreted as a *command* performed by the buyer to the world, plus an *assertion* performed by the world to the seller;
- the delivery can be modelled similarly to the payment.

3.1.4 Action and power

Fourth, activities are used as anchors for cognitive, motivational and institutional elements, informed by the illocutionary content of the messages. Our intuition is that *obligations* are prototypical motive for actions, while *power relations* are handled at epistemic level.

Therefore, in our representation, we refer in total to four layers, each of which addresses specific components:

- the signal layer — *acts*, *side-effects* and *failures* (e.g. time-out): outcomes of actions,

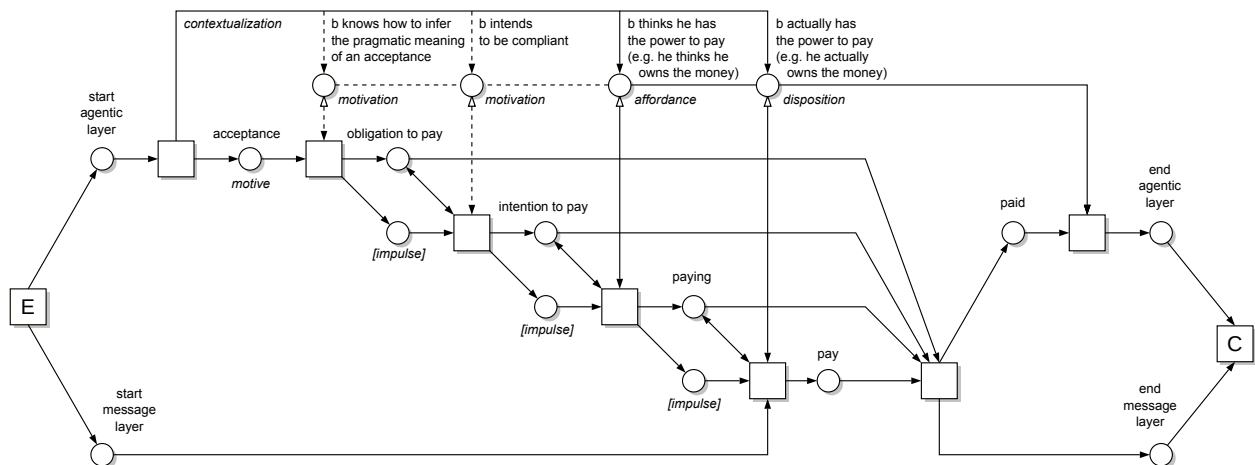


Figure 7: Full action pattern associated to payment.

- the action layer — *actions* (or *activities*): performances intended to bring about a certain result,
- the intentional layer — *intentions*: commitments to actions, or to nested intentions,
- the motivational layer — *motives*: events triggering the creation of intentions.

The last three layers compose the *agentic layer*. The closure of the sensing-acting cycle of the agent is guaranteed by the fact that certain signals, when perceived by agents, becomes motives for action. But how the motive is translated in intention(s)? How an intention is transformed in action(s)? What permits that an action effectively produces a certain result?

These questions can be answered introducing additional elements, dual to the previous components, such as:

- *dispositions*: contextual alignments of the agent with the environment (consisting of other agents and of the world actor) in respect to the actions he performs,
- *affordances*: perceived alignments of the agent with the environment in respect to his intents,¹⁰
- *motivations*: mental states catalyzing the creation of intentions.

Affordance and disposition Taking an intentional stance, all behaviours become intent-oriented. If

¹⁰In the literature—see for instance (Chemero, 2003)—the term *affordance* often refers both to the perceived relation animal/environment and to the effective (or dispositional) relation. In this paper we will associate the term *affordance* only to the first meaning, preferring the term *disposition* for the second.

the agent thinks that the environment *affords* his behaviour, he also thinks he has the power to achieve the goal associated to that behaviour.¹¹ Therefore, *affordance* practically corresponds to *perceived power*. Similarly, *disposition* is connected to *actual power*: it is a precondition to the consequences that a certain action of the agent will imply.

In the light of this analysis, we observe how these categories corresponds to the critical conditions reported on the MSC (e.g. *acceptability*, *ownership* in Fig. 4). The subjective evaluation of each of these conditions gives to the agent the *affordance* of the action resulting in that message. However, if the affordance is a sufficient condition for the performance of the action, it is only necessary for the intended outcome, where the contextual *disposition* plays a role. Considering our story, if the buyer starts buying, (it is like) he is assuming he has the power to do it—said equivalently, according to him the affordance of buying holds. Despite of his intent, however, the seller may be a fake seller, aiming to get the money without delivering. In this case, the disposition for a successful transaction would not hold, blocking the normal completion of the sale.

Motive and motivation In our framework, *motivation* refers to some mental condition that makes the agent sensitive to a certain fact, which becomes the

¹¹This assertion has no relation with consciousness. Intent can be implemented by *design*. A certain tool, or even a business process, is designed to achieve a certain goal in a certain environment, but artefacts or processes in themselves are not aware of such goal or environment. Intent and affordance may transcend actual performance, but they still exist.

motive for starting an action. As we observed before, obligations are prototypical reason for actions. Despite of that, not all obligations are followed by a performance. People comply with obligations when they have some motivation to: it may be for habit, convenience, respect to authority, or for fear of reinforcement actions. Motivations however often remain implicit in the story.

Visualization To sum up the concepts introduced so far, we have reported in the Petri net in Fig. 7 a possible reconstruction of the step of payment performed by the buyer. To reduce the visual burden, we have combined the layers *action*, *intentional*, *motivational* in one *agentic* layer. Nevertheless, it would be sufficient to add starting and finishing synchronization places for each layer to have a complete multi-layered model. The picture maintains in fact the vertical organization of our conceptualization: it is easy to recognize to which layer each element belongs.

The triggering motive is the acknowledgement of an acceptance (this could be the buyer's own acceptance to an offer, or the reception of a seller's acceptance to his own offer). The illocutionary content of the offer/acceptance entails the *duty* to pay of the buyer. This duty to pay is followed for instance if the buyer desires to be compliant. An obligation is then formed and used to construct the correspondent intention. The intention, if there is an action or a course of actions (i.e. *plan*) which is afforded by the environment (e.g. the buyer thinks he owns enough money), supports the selected performance. Finally, if the action is performed in a correct alignment (e.g. the buyer owns enough money), it results in the expected act.

In this type of Petri net we observe an interesting pattern — see Fig. 7: certain transitions are connected both to *impulse* and to *persistent* places. The first identify *events* (the occurrence of change), and the second conditions (the existence of a continuity).

3.2 Model validity

Each observed scenario can be explained by several interpretations: as the reader probably thought while reading our examples, there is not only one way of reconstructing the mechanisms that bring about the production of messages. The story provides only a *foreground*, which the interpreter can anchor to alternative *backgrounds*, explaining and completing it further. Further development of this research will be the study of a subsumption relation between scenarios (analysing their Petri nets) and consequently a distance operator which considers partial overlaps be-

tween them.

In this work, however, our focus is not the “scientific” validity of the systems integrated in the background, i.e. how far they represent what is the case in the experiential world; we accept that the choice of backgrounds is informed by the ontological commitments of the interpreter (e.g. judges giving alternative interpretations to a legal case). In our perspective, if an interpretation produces all the messages provided in the foreground, then, such model is *valid*. The process of selection becomes a matter of *justification*.

4 INTRA-AGENT SYNTHESIS

In the previous section, we have augmented the description of inter-agent interactions with intentional (and conversational) concepts. In this section, we will provide elements of the integration of our framework with current practices in MAS.

4.1 Agent-roles

Rationality is commonly defined as the ability of the agent to construct plans of actions to reach a goal, possibly referring to a hierarchical decomposition of tasks. In our case, agents do not deliberate, or better, the decisions they deliberate have been already taken.¹² Consequently, their behaviour is fully deterministic. Therefore, instead of considering a full account of agency, we opted to base our framework on a more constrained concept: the *agent-role*, which integrates the concepts of *narrative role* and *institutional role* in a intentional entity.

Agent-roles are self-other representations (Boer and van Engers, 2011a), i.e. used to interpret, plan or predict oneself's or others' behaviour. They are indexed by: (a) a set of *abilities*, (b) a set of *susceptibilities* to actions of others. In a social scenario, both abilities and susceptibilities become manifest as messages exchanged with agents playing complementary agent-roles. Thus, the topology of a scenario (e.g. Fig. 2) provides a fast identification of these indexes. Differently from objects (and actors), however, agent-roles are entities associated also to motivational and cognitive elements like desires, intents, plans, and beliefs.

¹²Part of the results of the deliberation are reported in the story, possibly together with elements of the deliberation process. Despite of these traces, however, the task decomposition occurs mostly in the mind of the modeler.

4.2 From scenarios to agents

A *scenario agent* is an intentional agent embodying an agent-role. It is deterministic in its behaviour: all variables are either set, or determined by messages from other agents. This determinism, consequent to the internal description of the agent, constrains the temporal order of messages between agents, but not the behaviour of an entire multi-agent system. In sections 2 and 3 we have constructed and used a message layer in order to preserve the synchronization given by the case. If we remove this layer, the system becomes non-deterministic. For instance, in a case in which there are two buyers and a seller, we do not know *a priori* who is the first buyer to accept.

Agents may communicate only with agents in their operating range, i.e. whose *message boxes* are reachable to them. Each agent may have different message boxes, and each of them can be epistemically associated to an identity.¹³

4.3 Computational implementation

For its structural determinism, the behaviour of an agent-role can be described in terms of rules, translating *logical* and *causal* dependencies. As we already observed analyzing Petri nets like in Fig. 7, some transitions are connected to places associated to *impulse* events and to places named after *continuing* conditions. This configuration easily relates our representation to *event-condition-action* (ECA) rules, commonly used in reactive systems. In respect to MAS theory, this connection has been exploited already in *AgentSpeak(L)* (Rao, 1996), a logic programming language for cognitive agents, extended and operationally implemented in the platform *Jason* (Bordini et al., 2007). The connection of *AgentSpeak(L)* with Petri nets has been extensively studied in (Behrens and Dix, 2008), with the purpose of performing MAS model checking using Petri nets. In the present work, however, we have a completely different objective: we started from a representation of the scenario on a MSC chart, we refined it with Petri nets patterns, and now we want to extract from this representation the correspondent agent-role descriptions (as agent programs).

4.3.1 Overview of AgentSpeak(L)/Jason

In order to proceed, this section presents very briefly two important constructs of *AgentSpeak(L)/Jason*. The first is the ECA rule associated to the activation

¹³Current MAS platforms refer instead to ontological identities.

of a goal. Put in words, with an imperative flavour: *in order to reach the goal, if certain conditions are satisfied, perform this plan of actions*. The equivalent code is something like:

```
+!goal : conditionA & ... & conditionZ  
      <- actionA; ...; actionZ.
```

Conditions represent what the agent thinks should be true, in that very moment, in order to be successful in executing his plan. The propositional content is written in a Prolog-like form. In addition, Jason permits the attachment of annotations (a sort of optional predicated parameters, expressed within squared brackets). As for actions, they are either direct operations with the environment, or !g (triggering the activation of a goal g), +b, -b (respectively adding and removing the belief b from the belief base).

The second construct we present is an ECA rule concerning the addition of a belief: *when a certain belief is added, if certain conditions are satisfied, perform this plan of actions*.

```
+belief : conditionA & ... & conditionZ  
      <- actionA; ...; actionZ.
```

Although similar, these rules have different semantics. The motivational component !goal disappears when the plan ends successfully and also when it fails (in this case, the event -!goal is triggered). The knowledge component belief is instead maintained.

4.3.2 Case example: buyer's payment

As example of synthesis, we are now able to translate the interpretation of payment illustrated in Fig. 7. This is an excerpt of the code of the *buyer* agent-role¹⁴:

```
+!accept(offer(Good, Amount)  
        [source(Seller)])  
      <- .send(Seller, tell,  
            accept(offer(Good, Amount)));  
        +obl(payload_to(Amount, Seller)).  
  
+obl(payload_to(Amount, Agent)) <-  
      !payload_to(Amount, Agent);  
      -obl(payload_to(Amount, Agent)).  
  
+!payload_to(Amount, Agent)  
  : owning(Money) & Money >= Amount  
  <- .send(w, achieve,  
        payload_to(Amount, Agent));  
      +paid_to(Amount, Agent).
```

¹⁴.send/3 is an action provided by the MAS platform that generates speech acts. The first parameter is the target agent, the second is the illocutionary force (tell for assertions, achieve for directives), the third is the propositional content.

For completeness, we included in the first rule also the generation of the speech act of acceptance. Neglecting this action, we have three rules, hierarchically dependent, and the last one performs the speech act for the payment. It is easy to observe a strong correspondence between these elements and the layers in Fig. 7. The first rule acknowledges the acceptance and generates the event/condition concerning the obligation (motivational layer), the second rule transforms the obligation of paying in intention of paying (motivational layer); the third rule checks the affordance related to the intent (intentional layer) and, if this evaluation is positive, performs the paying action (action layer). The action is then externalized to a communication module of the agent, interacting with the world/environment, which in turn will generate the actual consequence (signal layer).

To conclude, we observe that in this description the belief of having paid can be only partially aligned with the ontological reality. From the perspective of the agent, if the action has not failed, it is natural to think that it has been successful. In reality, however, something may block the correct transmission of the act to its beneficiary (e.g. a failure in the bank databases). The extent of such alignment is related to the focus of the feedback process checking the performance.

5 DISCUSSION

The modeling exercise running through the paper served as an example of operational application of a knowledge acquisition methodology targeting socio-institutional scenarios. Each representation we considered (MSC, topology, Petri net, AgentSpeak(L) code) has shown its weakness and strengths in this respect. Furthermore, the cross-relations between them are not simple isomorphisms. Despite of these difficulties, we think that using alternative visualizations is a way to achieve a more efficient elicitation (targeting also non-IT experts). In this line of thought, we plan to implement and assess an integrated environment for knowledge acquisition; the scalability of the methodology should be supported by the introduction of an adequate subsumption relation between stories, allowing faster elicitation of models.

From a higher-level perspective, the present work connects scenario-based (or case-based) modeling with multi-agent systems technologies. The idea at the base is that, in order to acquire representations of social behaviours, we need cases to be valid models, and we can validate them by their execution.

(Mueller, 2003) observes that, although several story understanding programs—starting from BORIS

(Charniak, 1972)—have used sort of multi-agent systems for their internal representation, this choice is not easy for the programmer: such agents are difficult to write, maintain, and extend, because of the many potential interactions. His experience matched with ours. However, we think that the connection of *agent-based modeling* with MAS is too strong and important to be easily discarded. As longer-term objective, we aim to couple on the same simulation framework designed systems (e.g. IT infrastructures) and representations of known social behaviours.

Scenario-based modeling MSCs (and collections of them, e.g. HMSCs) were standardized as support for the specification of telecommunication software, in order to capture system requirements and to collect them in meaningful wholes (Harel and Thiagarajan, 2004). Later on, other extensions, like LSCs (Damm and Harel, 2001) and CTPs (Roychoudhury and Thiagarajan, 2003), were introduced to support the automatic creation of executable specifications. The basic idea consists in collecting multiple *inter-object* interactions and synthesizing them in *intra-object* implementations. In principle, we share part of their approach. Our work promotes the idea of using MSCs, although integrated with intentional concepts. However, in their case, the target is a specific closed system (to be implemented), while in our case, a scenario describes an existing behavioural component of an open social system. At this point, we are satisfied by transforming the MSC of a single case in the correspondent agent-roles descriptions. The superposition of scenarios, with the purpose of associating them into the same agent-role, is an open research question.

Story understanding AI started investigating stories in the '70s, with the works of (Charniak, 1972), (Abelson and Schank, 1977), introducing concepts like *scripts*, *conceptual dependency* (Lytinen, 1992), *plot units* (Lehnert, 1981). The interest towards this subject diminished in the early '80s, leaking into other domains. (Mateas and Sengers, 1999) and others tried a refocus in the end of '90s, introducing the idea of Narrative Intelligence, but again, the main stream of AI research changed, apart from the works of Mueller (Mueller, 2004). All these authors, however, are mostly interested in story understanding. We are investigating instead the steps of construction of what they called *script* (Abelson and Schank, 1977). According to our perspective, common-sense is not constructed once, in a script-like knowledge, but emerges as a repeated pattern from several representations. Furthermore, we explicitly aim to take account of the integration of fault and non-compliant behaviours, increasing the “depth of field” of the representation.

Computational implementation Reproducing a system of interacting subsystems needs concurrency. Models of concurrent computation, like the Actor model (Hewitt et al., 1973), are implemented today in many development platforms. In our story-world, this solution would be perfect for objects. We would need instead to add intentional and institutional elements in order to implement agent-roles. The connection with another programming paradigm (intended to handle concurrency) will plausibly solve most of the problems of scalability that usually haunt MAS platforms, often developed by a logic-oriented community.¹⁵

Relevance The *Agile methodology for public administrations* (Boer and van Engers, 2011a) introduces the concept of agent-role, and targets the exploitation of agent-role descriptions, as components of a knowledge-base corresponding to the *deep model* (Chandrasekaran et al., 1989) of a social reality. Such a model can be used to feed design and diagnostic (Boer and van Engers, 2011b) applications with the purpose of supporting the activity of an administrative organization. The legal system in many areas presupposes the use of informal or semi-formalized models of human behaviour in order to operate. If we aim to support an administrative organization on those points, ABM is the most natural choice. In doing this, we recognize we are going in opposition to the current drift from *agent-based modeling* to *computational social science* (Conte and Paolucci, 2011). However, because of the generative aspect of the agent-role concept, our contribution is relevant to research in *behaviour oriented design* (Bryson, 2003) or similar, usually applied to robotics and AI in games.

REFERENCES

- Abelson, R. P. and Schank, R. C. (1977). *Scripts, plans, goals and understanding*. Lawrence Erlbaum Ass.
- Bal, M. (1997). *Narratology: Introduction to the Theory of Narrative*. University of Toronto Press, 2nd edition.
- Boer, A. and van Engers, T. (2011a). An agent-based legal knowledge acquisition methodology for agile public administration. In *Proc. of 13th Int. Conf. ICAIL 2011*, 171–180.
- Boer, A. and van Engers, T. (2011b). Diagnosis of Multi-Agent Systems and Its Application to Public Administration. In *Business Information Systems Workshops, Lecture Notes in Business Information Processing*, 97:258–269.
- Behrens, T. M. and Dix, J. (2008). Model checking multi-agent systems with logic based Petri nets. *Annals of Mathematics and Artificial Intelligence*, 51(2-4):81–121.
- Bordini, R. H., Hübner, J. F., and Wooldridge, M. (2007). *Programming multi-agent systems in AgentSpeak using Jason*. John Wiley & Sons Ltd.
- Bryson, J. (2003). The behavior-oriented design of modular agent intelligence. *Agent technologies, infrastructures, tools, and applications for e-services*, 61–76.
- Chandrasekaran, B., Smith, J. W., and Sticklen, J. (1989). 'Deep' models and their relation to diagnosis. *Artificial Intelligence in Medicine*.
- Charniak, E. (1972). Toward a model of children's story comprehension. Technical report, MIT.
- Chemero, A. (2003). An Outline of a Theory of Affordances. *Ecological Psychology*, 15(2):181–195.
- Conte, R. and Paolucci, M. (2011). On Agent Based Modelling and Computational Social Science. Available at SSRN 1876517.
- Damm, W. and Harel, D. (2001). LSCs: Breathing life into message sequence charts. In *Formal Methods in System Design*, 45–80.
- Diaz, A. F., Earle, C. B., and Fredlund, L.-A. k. (2012). eJason: an implementation of Jason in Erlang. *Proc. of 10th Workshop PROMAS 2012*.
- Fehling, R. (1993). A concept of hierarchical Petri nets with building blocks. *Advances in Petri Nets*, 674:148–168.
- Harel, D. and Thiagarajan, P. S. (2004). Message sequence charts. *UML for Real: Design of Embedded Real-time Systems*, 75–105.
- Hewitt, C., Bishop, P., and Steiger, R. (1973). A universal modular actor formalism for artificial intelligence. *Proc. of IJCAI 1973*, 235–245.
- Lehnert, W. G. (1981). Plot units and narrative summarization. *Cognitive Science*, 5(4):293–331.
- Lytinen, S. L. (1992). Conceptual dependency and its descendants. *Computers & Mathematics with Applications*, 23(2):51–73.
- Mateas, M. and Sengers, P. (1999). Narrative intelligence. In *AAAI Fall Symposium on Narrative Intelligence*.
- Mueller, E. T. (2003). Story understanding through multi-representation model construction. In *Text Meaning: Proc. of HLT-NAACL 2003 Workshop*, 46–53.
- Mueller, E. T. (2004). Understanding script-based stories using commonsense reasoning. *Cognitive Systems Research*, 5(4):307–340.
- Rao, A. S. (1996). AgentSpeak (L): BDI agents speak out in a logical computable language. In *Proc. of 7th European Workshop on Modelling Autonomous Agents in a Multi-Agent World*.
- Roychoudhury, A. and Thiagarajan, P. (2003). Communicating transaction processes. *Proc. of 3rd Int. Conf. on Application of Concurrency to System Design*, 157–166.
- Sileno, G., Boer, A., and Van Engers, T. (2012). Analysis of legal narratives: a conceptual framework. In *Proc. of 25th Int. Conf. JURIX 2012*, 143–146.

¹⁵Preliminary experiments of translating Jason on a functional programming paradigm, like for instance in Erlang (Diaz et al., 2012), have proven to be computationally very efficient.