

# Institutional Metaphors for Designing Large-Scale Distributed AI versus AI Techniques for Running Institutions

Alexander Boer and Giovanni Sileno

**Abstract** Artificial Intelligence (AI) started out with an ambition to reproduce the human mind, but, as the sheer scale of that ambition became manifest, it quickly retreated into either studying specialized intelligent behaviours, or proposing overarching architectural concepts for interfacing specialized intelligent behaviour components, conceived of as agents in a kind of organization. This agent-based modeling paradigm, in turn, proves to have interesting applications in understanding, simulating, and predicting the behaviour of social and legal structures on an aggregate level. For these reasons, this chapter examines a number of relevant cross-cutting concerns, conceptualizations, modeling problems and design challenges in large-scale distributed Artificial Intelligence, as well as in institutional systems, and identifies potential grounds for novel advances.

## 1 Introduction

These days, analogies carry easily between simulations of social settings and architectures for minds. We for instance casually speak of electronic and computational institutions as alternatives for traditional institutional arrangements, in recognition of the increasingly important role that the digital world, and automated decision making, plays within our social structures. Equally casually, normative multi-agent systems, based on institutionalist vocabulary, may be introduced as a design metaphor for large-scale distributed Artificial Intelligence (AI).

Drawing from existing work in agent-based modeling in the area of law, we examine in this chapter a number of relevant cross-cutting concerns, conceptualiza-

---

Alexander Boer  
KPMG, Amsterdam, the Netherlands. e-mail: [boer.alexander@kpmg.nl](mailto:boer.alexander@kpmg.nl)

Giovanni Sileno  
Complex Cyber Infrastructure, Informatics Institute, University of Amsterdam, Amsterdam, the Netherlands. e-mail: [g.sileno@uva.nl](mailto:g.sileno@uva.nl)

tions, modeling problems and design challenges that apply to both large-scale distributed Artificial Intelligence and agent-based modeling and simulation of social-institutional structures. Our aim is on the one hand to provide architectural indications to attempt to go beyond the contingent, task-centered, narrow view on AI; on the other hand to reflect on the continuity holding between institutional and computational domains.

### ***1.1 Modeling the Mind***

As a discipline, Artificial Intelligence (AI) started out with an ambition to reproduce the human mind as a monolithic computer program. The problem of reproducing intelligence in technology was taken on essentially as a *functional decomposition* exercise, with “intelligence” playing the role of the high-level function to be reproduced in a computer program through reconstruction from smaller functions realized by simple well-understood input-output modules.

The nature of the constituent primitive modules was already rather clear to early computer scientists: *logical inference* and *knowledge structures* should play a key role, because introspection, common sense psychology and current concepts of rationality suggested so (Newell and Simon, 1976). The *General Problem Solver* of Newell, Shaw, and Simon, (1959) may be considered the archetype of this approach, and is often used in this role in lectures about the history of AI. In architectures of this type, the problem of scaling up from the simplest toy examples of reasoning to plausible intelligent function, even in restricted cognitive niches, becomes one of managing a complex knowledge representation in a logic-based representation language, with the properties of the logical language enforcing some of the required modularity, or non-interaction, between the represented knowledge structures.

As soon as this required modularity breaks down, that is, when knowledge structures interact in unforeseen and undesirable ways, the construction of the higher level function — intelligence — may catastrophically fail, in an obvious way if we are lucky, for instance if the constraints enforced by the representation language are violated, or insidiously, through unintelligent behaviour, in many other cases. The AI system following this architectural paradigm, from an engineering point of view, in a sense lacks sufficient adaptive capacity, or resilience, to reliably deal with the addition of new knowledge structures.

### ***1.2 The Scale Problem***

The scale problem, in an architecture of this type, continually presents itself as a knowledge representation problem: there is always too little background knowledge available to correctly scope knowledge structures, and eventually the system will

make inferences that are not true, and, more importantly, obviously not intelligent. Hence we historically find attempts in AI:

- to uncover and study types of background assumptions — like for instance the infamous frame assumption (Morgenstern, 1996) — that are usually overlooked and threaten modularity;
- to codify massive amounts of common sense background knowledge for general AI use (Lenat, 1995), in the hope of eventually reaching critical mass;
- to find at least some knowledge structures that have universal appeal, and therefore exhibit the required modularity from an engineering point of view, and codify these as definitions into so-called top ontologies (Sowa, 1995) for reuse among the different modules of a distributed system, and
- to try to scope defeasibility in logical reasoning (Pollock, 1975) and restore the required modularity by designing new logics and representation languages that deal with defeasibility in a principled manner, containing undesirable interactions between knowledge structures.

Research that deals with these approaches is historically put under the *symbolic* AI tradition. Its main impact is perhaps in steering away the mainstream philosophy of this field from mathematical logic and a disdain of logical inconsistency and irrationality, and bringing it towards more psychologically valid views of reasoning and rationality.

Another, concurrent response to the scale problem has been to decompose the AI problem into two more manageable types of problems, licensing AI researchers to limit their attention to either:

1. studying well-defined specialized intelligent behaviours (e.g. playing chess, or recognizing human faces, or finding good solutions to a job-shop scheduling problem) in specific contexts of use, and
2. proposing overarching architectural concepts for interfacing specialized intelligent behaviour components. These are conceived of as agents, each competent in a limited number of intelligent behaviours, in a kind of virtual organization creating and maintaining the appearance of being *one* (more) generally intelligent agent to the outside world.

Solving the first type of problem has direct, viable commercial applications that justify financing research, but was in the past often belittled as not-really-AI due to its obvious lack of adaptive potential (funnily enough, the common-sense use of the term AI today refers typically to applications in this area).

Solving the second problem has traditionally received little attention, but in recent years is sometimes indirectly considered in the context of addressing the *responsibility* gap caused by increasing reliance on networks of autonomous systems (autonomous vehicles, autonomous weapons, ...).

### 1.3 Engineering a Mind as an Ecology of Minds

For the purposes of this chapter, we are mainly concerned with the second type of problem: architectural concepts. Archetypical for this response are the agents in the visionary *society of mind* of Minsky, (1988) (“The power of intelligence stems from our vast diversity, not from any single, perfect principle”, p. 308), but also the more concrete *intelligent creatures* of Brooks, (1991) in his famous “Intelligence without representation”. According to Brooks, the fundamental decomposition of intelligence should not be into modules which must interface with each other via a shared knowledge representation, but into independent *creatures*, that interface directly to an environment through perception and action, and exhibit a certain ecological validity within that environment, whatever that environment may be. The key takeaway here is to bring to the foreground the relationship between the creature and the environment, niche, or ecology it lives in.

Indeed, specialized intelligent behaviours in a specific niche can often be successfully isolated. Some problems, e.g. playing chess, may be attacked with specialized, explicit knowledge representation and an appropriate search algorithm. Others, e.g. face recognition, are less amenable to solution by introspection, and may be achieved not with techniques based on knowledge representation, but with machine learning algorithms on large databases of correctly labeled examples that exemplify the intelligent behaviour to be acquired. We are indeed increasingly becoming accustomed to the ability of computers to perform well in specific intelligent behaviours in specific contexts, often even beyond the levels of performance attainable to human beings.<sup>1</sup> At the same time, we are generally aware that these techniques do not generalize beyond a specific niche, and do not generally label them as intelligent, perhaps because we intuitively understand they lack the adaptive capacity that characterizes true intelligence.

The second problem of AI identified above, i.e. of finding architectural concepts for creating and maintaining the appearance of the system being *one* agent, has given us *multi-agent systems* (MAS), composed of multiple interacting agents in an environment, each with independent perception and action functions, and no global knowledge of the social system of which they are part. A key property of multi-agent system architecture is that there is no central controller agent with dictatorial powers over the others, as this would simply be the functional decomposition approach that does not work. Instead, more complex forms of organization and control are investigated, always with the ulterior goal of improving adaptive capacity of systems.

Although multi-agent system research has yielded many valuable theoretical insights into the functioning of organizations, no examples spring to mind of systems of practical value that successfully combine interesting sets of intelligent behaviours using this paradigm. It is however interesting to consider one well known system that does combine different intelligent behaviours and has demonstrated some success in finding the right *intelligent creature* for the right ecological niche: IBM’s Watson (Ferrucci et al., 2010). Watson has, amongst other applications, shown great

---

<sup>1</sup> See e.g. the AI index maintained at <https://aiindex.org>.

competence at winning the game show Jeopardy, and is famous mainly for that achievement. Clearly it does, from our perspective, only one thing very well: answering questions (QA). But the interesting thing is that it answers different kinds of questions using a number of completely different QA approaches with overlapping competence. Architecturally, it behaves like a coalition of intelligent QA creatures that *compete* against each other to produce the best answer to a question, and in the process of doing so it acquires feedback on its performance and becomes better at selecting the best competitor based on features of the question and the proposed answer. This competitive setting in itself is a convincing example of a winning departure from the traditional functional decomposition approach, and a move towards more interesting *organizational* metaphors.<sup>2</sup>

Multi-agent system technologies and concepts have thus far enjoyed most success indirectly as fuel for the *agent-based modeling* paradigm with interesting applications in understanding, simulating, and predicting the behaviour of real-world socio-legal structures on an aggregate level. However, agent-based modeling deals with a different type of problem, somehow dual to the ecological perspective illustrated above: settling which elements are required to specify one agent's behaviour. Whereas specifications based on simple reflex architectures or mathematical functions are plausibly sufficient for studying certain phenomena of scale that may be understood even using simplistic models of economic rationality, the human vocabulary concerning social actors often refers to intentional and normative categories, focusing primarily on qualitative rather than quantitative aspects of behaviour. Since the beginning of the 90s, research efforts have been put in place on *cognitive agents* and *normative multi-agent systems*, aiming to define agreed computational infrastructures building upon inferential interpretative templates such as the *theory of mind* and various theories of normative positions from law, institutional economics, and linguistics. Theory of mind is exemplified by the *belief-desire-intention* or BDI agent paradigm of i.a. Bratman, (1987), and Rao and Georgeff, (1995). The normative positions input usually draws from literature on norms (Alchourrón and Bulygin, 1971), normative relationships (Hohfeld, 1913), and *speech acts* (Austin, 1975). Typically, the main objective of these contributions is not to provide empirically plausible models of human reasoning, but to maintain a high-level but exact specification of the agent from an external perspective using terminology that remains comprehensible for humans.

## ***1.4 Purpose and Plan of the Chapter***

This short introduction gave an informal overview over general concerns shared by AI and computer science on the one hand, and cognitive, social, and legal sci-

---

<sup>2</sup> Examples of relevant competitive settings can be found in machine-learning methods too, see e.g. *generative adversarial networks* (GANs) (Goodfellow et al., 2014) or *ensemble methods* as *random decision forests* (Tin Kam Ho, 1995). However, the social constructs exploited in these solutions are still rather minimal.

ences on the other. At closer inspection, however, cross-pollinations between these domains are driven by different motivations. Limiting our view to a technical feasibility perspective, we are facing two different *complex adaptive system* design problems, to be solved with the same knowledge resources:

- the problem of modeling, simulating, and reasoning about *complex socio-institutional arrangements* like law, markets, business, etc.;
- the problem of designing *complex information system arrangements* based on design metaphors drawing from vocabulary on minds and institutions.

Application domains like electronic institutions are particularly fascinating for researchers from both traditions, because they make us face both design problems at once.

Our present aim is to introduce the main challenges at stake on both sides, in a way to inspire the reader to further reflection. For this purpose, we prefer a high-level, non-technical presentation over an exhaustive literature review. In the concluding section, we identify a number of potential grounds for novel advances in the application of institutional metaphors to distributed AI, and the application of AI techniques in modeling, analyzing, and evaluating socio-institutional arrangements.

## 2 Agency as Unity

### 2.1 *External and Internal Views*

In abstract terms, all problems reviewed in the introduction are expressions of a general tension between *unity* and *multiplicity*. Multiplicity attempts to recover where unity fails, or to achieve what cannot begin to be achieved by individuals alone. Multiplicity is made up of unities, which, seen as an aggregate, might qualify as a unity again. Recursively, unities might contain multiplicities in themselves. In this construction, two points of view co-exist: *external* (addressing a system-unity acting within an environment) and *internal* (addressing a system consisting of several components). An organization can be seen externally as one single socio-economic actor, consuming and producing certain resources (products, services, ...) for certain goals; internally, it can be seen as an arrangement of roles embodied by individual employees, but also by a network of business units, daughter organizations and coordinated partners. In the same way, a decision-support system or the guidance system of an autonomous vehicle can be seen as single pieces of software, producing results relevant to a given task, given certain computational and informational resources; or as wholes of interconnected modules providing functions necessary for the system to run. Not unexpectedly, internal and external views can be found in standard modeling practices, cf. UML; or in software engineering, with notions as *orchestration* and *choreography*. But why is something seen as a unity (externally) or as a multiplicity (internally)? More precisely, as we are dealing here with entities

meant to act within the world, what does it mean to appear to be one agent to the outside world?

## 2.2 *What Agents Look Like, Externally*

When interpreting a scene, observers frame it depending on the most pertinent cut determined by their point of view. Taking a famous example by Hart and Honoré (1985), a farmer may see a drought as producing a great famine in his country, but an international authority may instead put responsibility on the government of that country, because it has not prevented it by building up adequate reserves. Plausible criteria used to settle upon the stance for the observer to apply are the *informativity* (the capacity of drawing relevant conclusions with it— where relevance builds upon some value or derived interest for the observer), and the *cognitive effort* required (cf. *relevance theory* (Sperber and Wilson, 1986), *simplicity theory* (Dessalles, 2013)). Amongst the possible interpretative attitudes available to the observer, the *intentional stance* (Dennett, 1987) is more prone to prediction errors, but also the one producing highly informative results at reasonable cost. For instance, it would be difficult to interpret a child's action towards approaching some sweets by utilizing only information about his physiological state, or by evolutionary considerations. Similar considerations apply on the interpretation of the behaviour of groups, organizations, countries, or artificial intelligent systems as intentional agents.

## 2.3 *Intentional Ascriptions and Articulations as Reasons*

Taking an intentional stance, observers attribute to an agent beliefs, desires, and intents; and consider him as a *rational* entity, i.e. one which “will act to further its goals in the light of its beliefs” (Dennett, 1987), and they can construct, with an adequate approximation, what he will do next. Seeing this process as enabling the ascription of *reasons* to explain the conduct of an agent, we reveal the constituents of *mentalization* (Fonagy and Target, 1997), or more precisely, of *rationalization* of behaviour. A certain entity (a person, a country, a natural phenomenon) is seen as an agent because observers are *able* to ascribe certain beliefs, desires, and intents to it obtaining relatively correct predictions. The same applies when one agent observes introspectively his own behaviour, articulating his motivations as reasons for action.

Both *having* articulated beliefs, desires, and intents, and being *able* to ascribe beliefs, desires, and intents to an agent clearly assist in communication between social agents. Mercier and Sperber, (2011) make this point forcefully: evolutionary speaking, the main function of reasoning using knowledge structures *must be* argumentation.

However, a different outcome might be expected depending whether we take the agent or the observer point of view. Traditional experiments on *confirmation bias*

in human reasoning consistently show an interesting imbalance between producing and evaluating arguments: when arguing *for* a claim, humans seem to suffer strongly from confirmation bias (i.e. of selectively collecting or remembering evidence), but when arguing *against* the same claim they are capable of refuting weak, biased arguments based on selective use of evidence. Mercier and Sperber point out connections between the outcomes of these experiments and the rationality of the process of dialectical argumentation in a court setting, of which one could say that overall it aims (and generally succeeds) at producing truth even if the arguing participants are just trying to win an argument.

The court setting is ruled by the *burden of proof*, the duty to produce evidence for claims (Walton, 1988). Burdens of proof are systematically allocated to agents based on whether winning the claim improves one's position, and, when one considers legal norms, burdens of proof can generally be statically allocated to propositions in the condition of a rule. When considered from the perspective of burden of proof, the charge of irrationality in producing arguments loses its force. Instead, we wonder how rules become statically connected to communication settings, and which part of a rule exists to serve which interest.

## 2.4 Scripting Behaviour and the Resulting Ecology

But how it is possible for an observer to ascribe a certain behaviour to the agent, when in principle the number of courses of action the agent might be committed to is infinite? In the traditional approach to decision theory, a rational decision-maker decides a solution to a problem by maximizing over the candidate solutions. A similar optimization principle is at the base of the *homo oeconomicus* axioms used in classic economic theory: agents are self-interested, maximizing their utility functions. The initial infinite potential ascriptions are then reduced to a small list of (economically) rational ones.

There are many ways in which one can frame criticisms against those assumptions, and explain how we should deal with the boundedness of human rationality: decision problems in the real world are always potentially too complex, making it prohibitive and practically impossible to perform maximization over all possible alternatives. Heuristics, rules and procedures emerge as a natural response of the agent to unpredictable errors in selecting the best solution. Amongst the authors working on this subject, Heiner, (1983) identifies the root of this general phenomenon in the gap between an agent's competence at solving a problem and the difficulty of a decision problem in an environment, called the *C-D gap*. Heiner convincingly argues that for big C-D gaps an agent will generally perform more effectively and efficiently over time by following simple rules rather than attempting to maximize. A good chess player will for instance tend towards playing specific routine openings to capitalize on his experience in the types of games that develop from those openings, and get more creative in the middle and late game, even though standard interpretations of rationality would suggest that the best move becomes easier to

guess late in the game as complexity decreases. In social settings this tends to result in institutionalization of social interactions, as the actors involved all perform better if the interactions are predictable to the extent they can be. There are for instance a number of valid ways to conclude a sales transaction, but a supermarket will accommodate only one of them to structure interactions and reduce the cognitive load imposed on participants. The result is that agent engages in a kind of *ecological niche construction* (Bardone, 2011), associating relatively simple sets of rules to social roles he might play in interactions with an environment, dependent on what model he chooses for that environment.

## 2.5 Social Intelligence as Compressing Behaviour

Simplifying, an agent and an environment could be modeled as programs that exchange symbols, with the environment sending reward signals to the agent in reaction to its output signals. On similar lines, Legg and Hutter, (2007) define intelligent behaviour in terms of an agent's expected reward in an environment (and universal intelligence in terms of all conceivable environments). The agent has the task to predict which outputs will maximize the future reward. This general formulation covers for instance game playing (where one's competence is determined by the ability to predict the opponent's moves) and question-answering (predict the answer that will be rewarded). To be competent at prediction, the agent essentially has to guess the program that its environment (or opponent) runs. This is the same problem as compression of data: to discover a small program (the prediction mechanism) that reproduces the body of data (the symbols from the environment). This explains the establishment of the Hutter prize (2006): a reward for setting a new standard for compressing a large body of data: being able to compress well is closely related to acting intelligently. A similar intuition is formalized by works in *algorithmic information theory*, summarized in the expression "understanding is compression" (Chaitin, 2005).

As a concept of intelligence, *intelligence as compression* is a powerful abstraction, but also one that has clear limitations from an ecological point of view (Dowe, Hernández-Orallo, and Das, 2011). Firstly, it does not account for the origin of the reward mechanism. In lossy compression schemes (in audio and video compression) we for instance discard part of the data from the environment as noise rather than signal to be reproduced, because it does not matter to the human viewer (i.e. no reward is given). Without such an external reference, we however have no account of how our intelligent agent makes the distinction between useless and useful data, and always compressing everything is not an ecologically valid approach to dealing with the cognitive limitations of an intelligent agent in a complex environment. In other words, a theory of intelligence should be dealing with Heiner's C-D gap.

The second limitation deals with social environments, and is in this context obvious. The intelligence-as-compression concept accounts for learning by induction from examples, but not for social learning through interaction with other agents. It

cannot account for shared vocabulary between agents (i.e. ontology), not with argumentation, and not with the exchange of instructions and rules between agents. Any account of these omissions would need to explain how agents keep their programs sufficiently similar to exchange pieces of them with other agents.

Nevertheless, the core observation of the intelligence-as-compression approach remains of interest: Confronted with an environment, or with another agent, intelligence involves, among other things, the ability to select or create of a program that can be ascribed to that environment or agent, or to select or create a correlative program to drive rewarding interactions with that agent or environment.

## 2.6 Specifying Intentional Agents

The traditional AI perspective internalizes the intentional stance and considers agents as intelligent systems, entities performing certain actions in order to achieve certain goals, depending on their knowledge. Consider for instance the uncompromising external perspective towards agents taken by Newell in his seminal paper on the *knowledge level* (1982). Knowledge and goals are *ascribed* to an agent to explain its behaviour, and are separable from the structures and mechanisms that create that behaviour. In a sense, Newell proposes that the key problem of Artificial Intelligence is to create *models* of observed agent behaviour within the bounds of some design constraints, a model of deliberation loosely referred to as the *principle of rationality*. The question of how to reproduce mechanisms that create the behaviour is secondary, and in practice often easier if the design problem is well-defined, and the goals of the agent are easy to identify (as long as one is focusing on things like chess, or recognizing faces).

Considering goals as (a kind of) desires, this approach can be aligned to the traditional philosophical account of practical rationality, based on a *belief-desire* architecture (BD). Bratman, (1987) convincingly argued against such two-parameter characterizations, highlighting the role of *intentions* (I). In principle, a BDI framework allows us to consider that agents may have conflicting, inconsistent desires (e.g. to eat the chocolate and to follow a weight-loss plan), but those that are eventually selected for the actual conduct have to be consistent; it is this selection, reified as intention, that captures the *deliberative* aspect of agency.

However, for the reasons exposed above, in contrast to the traditional view of deliberative agent—according to which agents take rational decisions by weighing reasons—one could rather start from an *interpretative* basis, considering that agents might reproduce already established courses of action (or *scripts*). For their *ex-post* nature, scripts reify deliberations already occurred. Intentions, as commitments towards courses of actions, are then reduced again to low-order (in the sense of concrete, contingent) volitional states selected by higher-order (generic, structural) desires in accordance to the abilities and knowledge ascribed to the agent, and the perceived state of the environment.

Evidently, what is obtained through an interpretative standpoint can be remapped in generative terms, i.e., entering the agent's mind, the same mental states can be used as the drivers for his conduct. It is to *satisfy* his intents that the agent behaves in a certain way. It is unimportant to acknowledge whether such mental states actually exist or are epiphenomena of other unconscious activities of the brain. The principle of *responsibility* in our legal systems is based on the assumption that such mental states are significant to reason about human behaviour (cf. (Sileno, Saillenfest, and Dessalles, 2017)).

## 2.7 Agentive Positions

At further inspection, the *belief-desire-intention* BDI triad can be argued to be still incomplete. First, one can recognize a distinction between structural-abstract and contingent-concrete elements also at epistemic level. More importantly, the BDI template does not make explicit references to abilities, whose presence influence the action-selection process, nor to sensory capacities, determining which elements of the environment will be eventually perceived by the agent. To fill this gap, in previous works (Sileno, Boer, and Engers, 2015) (Sileno, 2016, Ch. 7) we recognized four primitive categories required to specify the behaviour of an agent: *commitments, expectations, affordances, susceptibilities* (CAES). Informally, they roughly correspond to what the agent wants, what he believes, what (he perceives) he is able to do, and what he is disposed to react to. Any instance of one of these classes is a *position* that the agent takes with respect to the environment.

**Commitments** (C), as a category, include all motivational states. There may be commitments which are *acceptable* for the agent (*desires*), commitments which are *preferred* by the agent, and finally commitments which are *selected* by the agent (objectives or *intents*), eventually triggering an *intention* associated to a certain course of action (or *plan*). The introduction of prioritization at the level of preferences serves to solve possible inconsistencies, allowing a conflict-free selection of the desires to which the agent may commit as intents. From a generative perspective, commitments can be interpreted as agentive dispositions having enough stability to drive action. In this sense, all living beings, even the most primitive, are functionally provided with the commitment to survive (as an individual, as a kin, as a species, etc.). On the other hand, we have also to take into account that not all reasons for action necessarily refer to concrete objectives: e.g. desires promoting *values* strongly influence the behaviour of agents, but at a deeper level in the rationalization.

**Expectations** (E) reflect the *situatedness* of the subject in the world. What the agent expects from the world is what he believes the world is, *actually* and *potentially*. The actual category mirrors the traditional definition of *beliefs*. The potential category identifies what, according to the agent, the world may bring about under certain conditions, and corresponds to the usual meaning associated with expectation, including causation aspects. For the intelligence as compression hypothesis,

expectations cover the role of program that the environment runs, used by the agent to maximize the available rewards for the given commitments.

**Affordances** (or abilities) (A) can be seen as an opportunities of action—possibilities of the agent to adopt certain behaviours, in certain conditions, to achieve certain results. Affordances interact with commitments to define which behaviour the agent will eventually select.

**Susceptibilities** (or sensory capacities) (S) are attitudes specifying the *reactive* aspect of agency; an agent is susceptible to a certain event if he exhibits some reaction to its occurrence, at the epistemic or motivational level.

If we posit a behavioural rule in terms of a commitment (C) to a condition to be satisfied, the minimal model triggering the action corresponds to the negation of the belief (E) that the satisfying condition does hold. This provides a simple basis to explain behaviour: the agent acts because there is a *conflict* between what he wants and what (he thinks) it holds. The CAES agent architecture allows us to capture additional functional dependencies, as for instance:

- *enactive perception* (from the set of active commitments Cs to a certain susceptibility S): the perceptual experience has to be entrenched with the course of action to which the agent is committed, mostly to be ready to anticipate potential (usually negative) outcomes by adequate monitoring;
- *affordance activation* (from the set of active expectations Es to an affordance A): because the epistemic conflict is specified at higher-level, an additional informational channel is needed to contextualize the capabilities of the agent to the current situation;
- *affordance inhibition* (from a commitment C', stronger than C, to A): a certain action may produce consequences which are in conflict with stronger desires; in this case, even if perceived, the associated affordance is inhibited to prevent the triggering of the action.

### 3 Institutions and Collective Agency

#### 3.1 Going Internal: Control vs Ecological Structures

The external point of view is basis for interpretation, the internal one is for *design*. In social systems the implementation layer consists of agents, therefore design corresponds to deciding structural arrangements that enable and support those agents to form and maintain robust organizations in order to achieve higher-order goals. The agents will operate through a mix of cooperation, coordination, and competition in problem solving, in all cases with limited information exchange and limited commitment to shared vocabulary and rules.

Evidently, by distributing the (operational, computational) charge to networks of independent agents, we are naturally evolving from *control* structures to *ecological* structures. The difference between the two categories is best visualized by the

distinction (originating with Deleuze and Guattari, (1980), recently revisited by De Landa, (2006)) of the *totality* vs *assemblage* conceptualizations of wholes, summarized in the following table:

<i>Totality</i>	<i>Assemblage</i>
organic composition (e.g. heart in body)	ecological coupling (e.g. symbiosis)
components specified by relations of <i>interiority</i> : all their properties are manifest	components specified by relations of <i>exteriority</i> : only part of their properties is manifest
components exist only as part of the system	components exist even in absence of system
dependencies logically necessary	dependencies contextually obligatory
failures compromise the system	failures irritate the system

The passage from totality to assemblage requires several conceptual steps. For the sake of the argument, let us imagine we start from a monolithic software implementing an IT service, made up of several internal modules. The first step is adding redundancy in the system, i.e. components potentially providing similar functions. If there is only one module that implements a required function for the IT service to run and this module fails, the application will stop working properly. If there are a number of modules that might be invoked providing overlapping functions, the service resilience will generally increase and the problem will become rather one of an economic nature—i.e. of settling on an adequate resource distribution strategy amongst these modules. However, the IT service example still considers a core module that is dispatching resources to the others, depending on its requirements (just like MapReduce (Dean and Ghemawat, 2008) for service resilience relies on a scheduler allocating resources). As a second step this constraint needs to be relaxed. A pure assemblage has no core module: there is no co-coordinator putting preferential requirements in the foreground. (In nature, maintenance functions emerge as a *selection effect*: assemblages not able to sustain themselves disappear.) Third, components of an assemblage exist independently of whether the assemblage holds or not. In other words, components of a totality are fully specified by their actual properties within the system (*relations of interiority*), whereas components of an assemblage have dispositions which are not amongst those manifest in the context of that specific assemblage (*relations of exteriority*).

Considering the components of the assemblage as agents, whose intent satisfaction counts as reward obtained by the environment, we can draw the following general templates:

- *competition*: agents are committed to the same target, but its satisfaction by one produces a (physical or symbolic) scarcity for the others;

- *cooperation*: agents have dependent commitments, i.e. the satisfaction of one agent's intent is required for or facilitates the satisfaction of another agent's intent; symmetric dependencies are at the base of mutualism and enable maintenance functions (as in the symbiosis example);
- *coordination*: a structured mechanism (resulting from an individual or collective agency) distributes rewards and punishment to the agents specifically to obtain higher-order goals.

A discriminating factor between cooperation and coordination schemes can be the presence of explicit signalization. For instance, symbiosis exists even without communications.

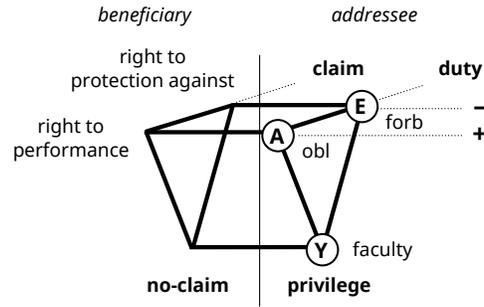
### 3.2 *The Central Role of Failures*

Institutions are a prototypical example of mechanisms of coordination: they build upon symbolic means, and form an infrastructure providing rewards and punishments (the anecdotal “carrots” and “sticks”) to social actors, modifying those available from the non-institutional environment. Note how competition aspects are in general mostly extra-institutional (e.g. the choice of a sale price in a market). The practical function of the legal system, as an institution, is to intervene when an (institutional) failure supposedly occurs in social interactions, i.e. when the institutional expectations of one of the parties involved are not met. *Ex-post* judicial interpretations are meant to make explicit the normative positions of the parties before the failure, and then to associate the institutional response if the failure is confirmed, building upon the sources of law. In this frame, normative sources are then used as reasons to enrich behavioural models ascribed to the parties.

### 3.3 *Hohfeldian Prisms*

Hohfeld's analysis of fundamental legal concepts (Hohfeld, 1913) starts from a similar interpretative consideration and captures two distinct dimensions—the *obligative* (or *deontic*) dimension and the *potestative* (or *capacitive*) dimension—of the legal relations holding between two social actors that are bound by legal provisions. The resulting framework brings two crucial innovations.

First, the consideration that concrete normative relations cannot be expressed on the mere basis of deontic modalities. For operational purposes, it is crucial to make explicit which party is the *addressee* of the normative proposition and which party is the *beneficiary*, i.e. whose interests are protected or promoted (cf. the teleological aspect highlighted by Sartor, (2006)). Thus, the notion of *claim* embodies the idea of right as the protection of an interest via a corresponding *duty*. A *privilege* corresponds instead to the absence of duty, and when it holds the other party has *no-claim* to advance.



**Fig. 1** First Hohfeldian prism, containing obligative positions.

By observing the intuitive correspondence of duty with obligation and privilege with permission (in the common meaning of *faculty*, not the usual formal meaning), we added a negative dimension to the traditional Hohfeldian square, obtaining the first Hohfeldian prism (Fig. 1, where A, E, and Y are the positions of the deontic triangle of contrariety) (Sileno, Boer, and Engers, 2014).

The second innovation is the explicit consideration of the dimension of institutional change, centered around the notion of *power*. Hohfeld insists also on the characterization of institutional power with volitional control, that is, with intentionality. A person holding such a power has the institutional *ability* to deliberately alter legal relations (e.g. transfer of ownership) by performing certain acts.

Rather than using terms of addressee and beneficiary when considering the action, the two parties can be distinguished as (potential) *performer* (enacting the power) and *recipient* (suffering from the enactment). Correlative to power, *liability* means being subjected to that power, while the opposite *immunity* means to be kept institutionally untouched by the other party performing the action (who, in turn, is a position of *disability*). As before, introducing a negative dimension we unveil the second Hohfeldian prism (Fig. 2), discovering the neglected positions of *negative liability* and *negative power*, relevant to undermine institutions (for an extended analysis of the Hohfeldian prisms, see (Sileno, 2016, Ch. 4)).

The visual representation given by the prisms (or squares, as in the original contribution) makes explicit the *symmetry* and *duality* of the relations between two parties. Focusing on a specific perspective (e.g. that of the addressee), two *positions* (or three, counting the negative attitudes) are available to describe in which situation that party stands with respect to another party. This position is strictly linked to the position of the other party. One may think of those diagrams as a game board in which when one player moves, he moves both checkers at the same time (*correlativity axiom*). Thus, the difference between, for example, duty and claim is just one of point of view, as they describe the same binding.

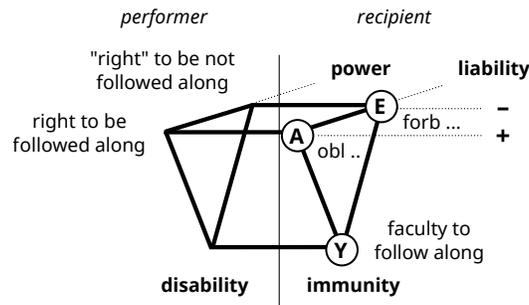


Fig. 2 Second Hohfeldian prism, containing potestative positions.

### 3.4 Interface between Normative and Agentive Positions

At a theoretical level, the four agentive categories of commitment, expectation, affordance and susceptibility (cf. CAES framework) can be put in direct correspondence with the four normative categories duty, claim, power and liability, interpreting the environment as the correlative party to the agent (Sileno, Boer, and Engers, 2015). The distinction can be traced as one of *intrinsic vs extrinsic* attitudes of one agent's behaviour. Social norms (including legal norms) provide reasons for the agent—by attempting to influence the rewards of the environment—to promote or demote certain action-selections in his conduct (via obligations and prohibitions), or even create the possibility of certain action-selections (via institutional power).

Let us divide all institutional knowledge in class-norm (N) and token-fact (F) types of components. Modifying rewards of the environment, both norms and institutional facts supposedly play a role in the expectation category within one agent's reasoning. Intuitively, to facilitate normative alignment between agents, norms should be as much as possible accessible, and therefore require a public medium of transmission. The case of institutional facts is a bit different; to protect sensible information, they should be in principle shared only on a institutional coordination basis. However, this distinction is not so strong as it seems: the universal ambition of the modern rule of law can be seen rather as a special case, consequence of the large base of addressees; companies do not necessarily want to share their policies widely, and social groups may have unique customs not explained to outsiders, with the effect of facilitating the distinction of in-group and out-of-group interactions.

Adding these institutional elements, we can make explicit several types of *transformational* powers at stake:

- *directive power*, of bringing the commitment C to action;
- *operational power*, of performing the action to achieve a certain goal, by means of the ability A;
- *enabling power*, of activating the ability A (cf. the dual *disabling power*);
- *regulatory power*, of transforming the individual commitment C to the norm (in the sense of collective commitment) N;

- *publishing power*, of reifying the norm N on an artifact publicly accessible;
- *interpretative power*, of interpreting observations, producing beliefs and expectations E;
- *monitoring power*, of perceiving a certain input, by means of a susceptibility S;
- *attention power*, of activating the susceptibility S (cf. the dual *diverting power*);
- *declarative power*, of transforming the individual belief E to an institutional (collective) fact F;
- *registering power*, of reifying the fact F on an artifact.

### 3.5 *Distributed Agency*

Hypothesizing that the previously identified functions are correct, we should find them independently from whether the agency is implemented in a control-based or an ecological system, natural or artificial. In effect, one can recognize in the previous section a basis for the famous *trias politica*. The *executive*, *legislative* and *judiciary* branches can be seen as concentrating on themselves the directive, regulative and interpretative powers of a collective agency. The principle of *separation of powers* can be explained as aiming to maintain the balances present in the standard reasoning architecture: N puts constraints on E, but E has the power to re-contextualize them by producing F, in turn inhibiting/enabling the abilities A available for C. Finally, even if C could be in principle transformed in N, N contains also previous decisions, taken by supposedly different power-holders, plausibly establishing constraints on C's regulative maneuvers. If the same actor holds, e.g., both the directive and interpretative powers, there is a short-circuit in the architecture because interpretation could be forced to enable this actor's commitments, rather than respecting those consolidated in the norms.

Most interventions of the judiciary occur on an asynchronous basis: the triggering event consists of a social actor (the plaintiff) that supposedly experienced an institutional failure and goes to the court. This means that the monitoring power is left to each social participants.<sup>3</sup>

The increasing presence of global online platforms centralizing socio-economic transactions, as well as the agglomeration of institutional administrative repositories is also opening possibilities of upstream monitoring by the entities having access to such information. De Mulder and Meijer, (2012) have argued for an extension of the principle of separation of powers to a *tetras politica*, including these potential monitoring branches. According to our architectural analysis, their argument is sound and it should even be enriched also by focusing on the actors maintaining registering power, probably the next frontier of innovation if distributed ledgers or related technologies will eventually enter into the institutional operational cycle.

---

<sup>3</sup> Media outlets are instead characterized as actors that (generally compete to) hold attention power.

### 3.6 Maintenance Functions

Although we have recognized in collective agencies functions ascribed to a single agent, we have not explained how this is actually possible. For systemic maintenance, we expect the existence of some *enforcement* mechanism, that is able to guide and maintain to a good extent individuals in principle independent to form an assemblage. (Consider again the symbiosis example, the contextual obligation of interaction is a consequence of the positive reward of such an organization with respect to absence of interaction.)

Enforcement actions generalize this principle, as they supposedly provide an institutional response to failures and strengthen the general compliance using *rewards* and *sanctions*, but the enforcement is not necessarily due to mere legal mechanisms, but also to social or physical dispositions (e.g. regulations of right-hand/left-hand traffic are issued in the presence of a sort of *natural* enforcement: not complying with the correct driving direction would have direct and possibly terrible consequences). However, legal transactions have the crucial advantage of putting in the foreground explicit protections to the risks of lack of coordination between the parties (see e.g. the architectural view of the problem presented in (Sileno, Boer, and Engers, 2020)).

The general evolution of contemporary enforcement practices—observable not only in modern legal systems, but also in school, parenting and informal groupings—has been nicely synthesized by Dari-Mattiacci and Geest, (2013) as “the rise of carrots and the decline of sticks”. The authors propose a theory explaining this tendency. In short, from an economic perspective, punishment-based systems are more efficient in simple settings, i.e. when the burden of compliance can be well specified and distributed amongst the involved parties. Reward-based systems, on the other hand, are more efficient when the regulator has difficulties specifying the burden, or the burden is not well distributed amongst the social participants. As modern societies become more complex, heterogeneity of contexts and tasks is increasing, and so reward-based systems are becoming more common.

At first glance, we could write any contract either in carrot-style or in stick-style and they would be identical, at least from an algebraic point of view. However, experiments conducted in behavioral economics on similar formulations found out that the two rephrasings are not the same (Kahneman and Tversky, 1979). Even at the analytical level, Dari-Mattiacci and Geest observe that there is an intrinsic difference due to the *monitoring* component: if monitoring is not perfect, the two formulations are different because a party that is not monitored in a carrot-based system counts as being non-compliant; whereas, in a stick-based system, it counts as being compliant. In practice, as observed by Boer, (2014), a reward regime assumes that most people are non-compliant; a punishment regime assumes that most people are compliant. Boer proposes then the following *evidence criteria* to decide whether an enforcement system is based on carrots or on sticks:

- a *reward regime* requires the production of evidence of compliance,
- a *punishment regime* requires the production of evidence of non-compliance.

These criteria switch the attention from internal, subjective aspects relative to the parties to external aspects specified in regulations. Interestingly, these are related to the *burden of proof*, usually distributed amongst claimants (for evidence of non-compliance) and addressees (for evidence of compliance).

## 4 Pluralism as Ecological Diversity

### 4.1 *Epistemological Pluralism*

Interestingly, the dichotomy totality vs assemblage captures also the distinction between formal and informal conceptual systems. The difficulty of constraining natural language in formal semantics is well known. The regularity of association of linguistic terms to a certain interpretation is a consequence of a relative regularity of the environment, of the perceptual apparatus, of the knowledge and of the commitments of the locutors—but a change within one of those aspects might entail a change in the associations as well. Taking into account this phenomenon means to accept *epistemological pluralism*.

Furthermore, beyond *inter-agent* pluralism (agents, for being independent, do not share the same knowledge, nor the same interests), there is also an *intra-agent* pluralism to be considered, due to *ontological stratification* (Boer, 2009). For instance, a road may be seen as a curved line while looking at directions on a map, as a surface while driving, or as a volume while putting asphalt to build it. The alternative cuts, as with the intentional stance, are due to the different properties of the entity which are relevant to the task in focus. Any alignment of vocabularies or ontologies cannot be successful without an adequate analysis of the underlying commitments associated to the task domain.

Normative systems, which are defined symbolically, prototypically suffer from this problem, because the possible contexts in which they might operate cannot be predicted in advance, but at the same time, they can also be seen as a prototypical solution to the problem, because they attempt to consolidate goals in higher-order, abstract form, leaving the specification of further contextualization at need.

### 4.2 *Normative Pluralism*

But there is another level of assemblage to be taken into account for normativity: agents belong to a single normative system only in a ideal case. Think of all the rules and norms, written and unwritten, that we encounter in daily life. At home we did certain chores because we at some point agreed to a distribution of them, or perhaps simply because they are expected of us. We followed both traffic rules, and informal driving etiquette, on the roads while traveling to and from work. We

parked the car following rules and directions, greeted colleagues following social conventions, and perhaps followed directions about where not to smoke or not to drink our coffee and where to discard our paper coffee cup. Finally we are writing this chapter following grammar and spelling rules, instructions about structuring chapters, and conventions about citation and plagiarism, working to meet a deadline. These norms belong to different types, and we recognize them as part of different complex systems of norms. Some are based in law, others in social conventions of society at large, in social conventions of the academic field, those of a particular building, or even of a small private circle. Any of these norms we may have noticed or not noticed, regarded or disregarded, and followed or violated, consciously or unconsciously. And for the most part we hardly think about them.

Normative pluralism does pose two really interesting questions (see e.g. Twining, (2010)). The first question is: *how do we notice norms, and detect and resolve conflicts between pairs of norms as they arise?* When two norms are in conflict they are so because they either:

- disaffirm each other, i.e. appear to label the same intention as both strongly permitted and not permitted, or
- appear to require of us that we adopt two intentions that jointly are inconsistent with our beliefs and hence cannot be executed jointly if our beliefs are correct, for instance if we ought to have our cake and eat it.

From the point of view of the norms we might state that norms may be grouped into a normative order, which is an institutionalized system of norms aimed at ordering social relationships between members of a community or group. Conflicts between these norms are largely ironed out within the system of norms, and norms for determining applicability and priority of norms may determine the outcome of new conflicts. This is however only possible to the extent that agreed-upon norms on conflict resolution and prioritization can be established within the community or group, noticing applicability of norms is not itself the overriding concern, and the outcomes of the conflict resolution process remain acceptable to the members of the community, i.e. the right norms are applied.

### ***4.3 Value Pluralism***

The second question issued from normative pluralism is: *how do we measure how norms perform overall if they are to be effective in a plurality of norm environments?* Systems of norms range from the public, with universal appeal, to the most private. Thus, the problem of measuring norm performance is relevant at all levels of granularity, and from an intra-agent and an inter-agent perspectives. Individual agents, organizations, and whole communities form judgments about the performance of individual norms, and the problem of forming opinions about individual norms in a plurality of norm environments is rather ill-defined if it is the norm environment itself that frames what rational action is.

One might naively think of this problem as a quantitative one, of measuring overall levels of compliance. From the perspective of the values served by norms this is however typically ill-advised. A norm performs well if it works, or is complied with, in those environments where it is useful, and does not work, or is violated, in those environments where its value is outweighed by more important values served by other, competing norms that (should) take precedence. As such, overall compliance levels do not necessarily mean something, if they are not properly scoped to a niche in which norms should be complied with.

#### ***4.4 Pluralisms and Social Niches***

A general pattern could be traced at this point. Intra-agent epistemological pluralism emerges from simplifying the agent-environment coupling relative to functional niches, determining the explicit articulation of several agent-environment couplings. Inter-agent pluralism aggregates those articulations, and the commitments exploiting them, resulting in a set of *agent-roles* available within a social niche. Normative and value pluralism can be seen as components of the previous ones, aggregating respectively *social coordination* (institutional) and *reward model* aspects. Other possible inflections are *political pluralism*, necessary base for democracy; and *legal pluralism*, a special case of normative pluralism, necessary requirement for jurisdiction in international matters.

In other words, pluralisms come together with accepting an ecological view of social systems: they mirror the *diversity* (of aspects) of the social niches available within a society, at a certain point in time. This implies that rules can be understood and modeled only within the context of social roles (in turn associated to niches). All effort to capture regularities, rules or associated epiphenomena without considering the niche that is using them is going to fail, because as soon as the relative distribution of niches in society changes, the model foundations will change too. This observation is compatible to accounts highlighting the *material* dimension of law, according to which the law is primarily constituted and performed out of dedicated practices pertaining to a certain social assemblage (Philippopoulos-Mihalopoulos, 2014; Pottage, 2012). Vice-versa, from a system-design perspective, this view entails that all artifact capturing norms (norm as in *normativity* but also as in normality) need to be considered in an ecologically sound computational architecture, able to assimilate or accommodate adequately to any environmental configuration. This requirement suggests the introduction of a *normware* level of conception of artificial devices, beyond hardware and software (Sileno, Boer, and Engers, 2018).

## 5 Discussion

### 5.1 *Institutional Metaphors for Large-Scale Distributed AI*

With the increasing use of distributed, computational, possibly autonomous, systems, of technologies based on distributed ledgers and smart contracts, of the Internet of Things, etc., it is relevant to apply the previous conceptualizations to assess to what extent the institutional mechanisms identified in the legal domain have a correspondence in computational social systems.

Despite the name, “smart contracts” (see e.g. (De Filippi, Wray, and Sileno, 2020)) do not embody any problem-solving method, nor are specified by assigning normative positions to parties as in usual contracts. Their main innovation, as with the block-chain, is the *distributed ledger*, used both for “contract” publication and for registration of the related transactions, removing the requirement of an explicit maintainer (e.g. a bank, a public administration, a notary, etc.). They are creating the basis for a potential infrastructural centralization of registration power. Unfortunately, by collapsing normative functions to the implementation layer, these artifacts are fundamentally opaque to users. Second, they do not enable architecturally the negative feedback of interpretative power on directive power for novel or exceptional contexts not taken into account at design time. This heavily undermines the reasonableness of the solution for institutional operationalizations. In spirit, they do not differ from *paperclip maximizers* (Bostrom, 2003).

In recent years, many efforts have been directed towards the development of *secure* operating systems. Traditionally, most implementations build upon solutions in the spirit of *access control lists* (ACL), i.e. mapping each user and object to a series of operations allowed to perform on it (e.g. read, write, execute) (Ferraiolo, Barkley, and Kuhn, 1992). These permissions are usually called also privileges, or authorizations; but at closer inspection, it is a conflated version of the homonym Hohfeldian position. Without such a permission, the user is *disabled* to perform that action, not prohibited. For their dematerialized content, computers are not so different from institutional systems: they both build upon symbol-processing mechanisms. In this sense, writing a file is not a physical operation, but an institutional operation, and so, to perform it the user is required to have the correspondent power. Capability-based security models are implicitly based on this, using communicable tokens for authorization known as *capabilities* (e.g. Levy, 1984). The *principle of least privilege* (or least authorization) requires that capabilities are assigned only within the actual purpose of the application. However, considering authorization merely as power carries additional concerns. In actual socio-legal settings, the most common usage of permission of *A* is when the agent has the (usually physical) power to perform *A*, but is prohibited from doing it. More in general, permission is needed because power has a too low granularity to capture higher-order effects. For instance, even if a single action *per se* may be allowed (enabled), that action in a course of actions, or together with other actions, may bring about unauthorized

results (consider e.g. *denial of service* (DoS) types of attacks exploiting trusted applications).<sup>4</sup>

Evidently, failure cases extracted from *ex-post* evaluations can be used effectively for monitoring potential preparations of known schemes, but the actual problem, resolved in human societies by using deontic directives, is to specify principled reasons as a basis or anchor to qualify *new* types of failures within the existing normative framework.

Focusing now on AI methods, the most interesting results obtained in these last years comes from an hybrid application of deep learning methods with reinforcement learning (starting from AlphaGo, (Silver et al., 2016), ). Also genetic algorithms can be interpreted as issued through evolutionary reinforcement. For the nature of the tasks, most reinforcements can be associated to a centralized reward system, providing the agent with something positive/negative (that he would not have had otherwise) if he attains an outcome qualified positively/negatively. These are only two amongst the six (primitive) possible figures of reward/punishment regimes that can be identified in institutional systems (see (Sileno, 2016, section 9.5)). But there is a more profound problem that undermines their generalization towards less specialized problems: *the need of specifying clear-cut reward/punishment* functions.

In effect, AI researchers and practitioners (and in general problem-solvers) tend to think that the identification of goals is the easy part—and how to get to these targets the hard part. However, when conceiving applications that are required to adapt to the user, or to the social environment, this presumption rapidly collapses, and could not be otherwise for at least three reasons: for the many possible configurations of preferences between social participants, for such configurations being highly contextual, and for most preferences to be tacit.

The main weakness of contemporary AI lies in trying to capitalize too much on optimizing and reasoning about the evidence and options for action (*generate & test* paradigm) within known requirements, constraints and problem formulation, rather than looking into underlying phenomena of niche construction and adaptiveness, and finding requirements as related to social roles. Deciding to target social roles, rather than formulating requirements on an individual basis, follows from Heiner's theory and is cognitively supported by the observation that also humans actually perform better at deontic reasoning than evidential reasoning (Mercier and Sperber, 2011), and acquire this ability earlier (as small children).

Interestingly, the AI & Law discipline (Bench-Capon et al., 2012) can be thought as originating from reversing the general AI attitude, that is, by focusing strongly on the conflict between goals, and the mechanisms by which we acquire and select them (*normativity*, or, even more deeply, *values*). Capitalizing on this experience, even considering the automatic discovery of ecological niches too ambitious, one could still aim for the automated design of compositions of social roles meant to achieve

---

<sup>4</sup> As an additional example of engineering “blindness” to social semantics, consider the endless discussions about of what fragment is logic to be offered for knowledge representation in the semantic web, mostly focusing around computability and halting problem vs. the fact that, as soon as you open up an interface to the outside world through a SPARQL interface, the system can be trivially put out of service by a DoS attack regardless of the fragment of logic adopted.

certain given requirements. The core problem in obtaining such construction would be of detecting and solving the internal role conflicts, and doing so by internalizing pluralism (normative, value, epistemological). A critical part in the design would lie in settling adequate data and knowledge roles, the ones which carry data and knowledge between behavioural niches. An improvident centralization of data, for instance, might enable the utilization of the collected information for purposes non-intended by the data-providing agents. This unveils the need of a *principle of data-information minimization*, i.e. that the provision/collection of data is calibrated to the adequateness of reasons for access to information (dependent on the social role attributed to the requestor and thus the social context), rather than a token-based principle of least privilege (dependent on the requestor). In other words, it urges for a revisit of the function of responsibility attributed to computational entities (see e.g. Sileno et al., 2020).

## 5.2 AI Techniques for Running Institutions

In the current globalization trend, all socio-economic actors (private individuals, private companies, NGOs, and public institutions) become more and more active at an international scale, and institutional interdependencies pass from exceptional to normal operative conditions. Operating concurrently within different legal jurisdictions, these actors are subject to an ever-changing compound of rules (e.g. national provisions, international treaties, standards), whose complexity naturally increases with the number of international activities maintained. This carries risks of misalignment, which may lead such actors to suffer enforcement actions and institutions to suffer failures. On the other hand, as their activities occur at international level, those actors have the possibility to establish practices—for instance, aiming to reduce the burden of compliance with environmental, labor, or privacy regulations; or to implement schemes of tax avoidance, tax evasion, or money laundering—which are not yet (or impossible to be) recognized by formal institutions at national level.

From a theoretical perspective, increase of complexity means an increase of the C-D gap for the agents, and therefore a pressure towards establishing simpler, more manageable operational rules, at the potential expense of the normative elaboration conducted upfront and of the diversity of cases that may occur, and that would instead require, to serve the citizens' interests, a dedicated response. In this context, AI, and in particular AI & Law, could provide methodological and computational support for consolidating the *rational* functioning of collective agencies.

Beyond classic problems-solving tasks for operational purposes (e.g. *planning*, *scheduling*), here we refer in particular to methods supporting higher-order functions of rational agency as the one highlighted above. In current practices, one can observe little attention to pluralisms within organizational settings, caused by focusing on implementation and optimization of “normal” provision patterns—the so-called *happy flow*. Consequently, even if presumption of compliance utilized at the rule design phase is generally untested, patterns of failure are put outside the

scope of attention. Thus, even if numerically negligible, failures in service delivery absorb proportionally more and more resources, and result in a more complicated experience for the service consumer (and so the anecdotal problems of *machine bureaucracy*, see for instance (Seddon, 2008)).

Symptomatic of this situation are unrealistic ways to quantify and measure performance, little attention to the *systemic* meaning of failures, and no implementation of a systematic diagnosis process that would reduce the operational burden (see e.g. (Sileno, Boer, and Engers, 2017)). On a par with this, there are excessive expectations about the information that can be obtained from big data analysis, which are implausible because the niches in which data is generated are completely neglected.

On the other hand, collective agencies risk to take an attitude of *paralyzing realism*, when they put excessive attention on capturing pluralism in society matters at high granularity, as they become unable of completing the aggregating requirements phase necessary to proceed to the directive and then operational or regulative phase.

Both machine-bureaucracy and paralyzing-realism scenarios results from not settling upon an adequate systematization of environmental modeling (too little, or too much). The solution lies in establishing adequate networks in which to distribute the burden of capturing social requirements and scenarios of non-compliance, to be transmitted, maintained and processed for support to the design/development phase and of the operational phases (e.g. with automatic model-based interpretation, including diagnostics, see (Boer and Engers, 2011a,b; Sileno, 2016; Sileno, Boer, and Engers, 2017)).

### 5.3 *Converging Opportunities*

Despite current hopes, even introducing *big data* analytics methods, the qualitative depth of the knowledge available to decision-makers of collective agencies cannot substantially improve. Such value can be captured only by the introduction of more principled modeling and simulation methods targeting *social niches*, together with knowledge maintenance, refinement and alignment techniques. This choice would have as a direct consequence a concrete reduction of the C-D gap, and then would provide a strong support for the design of more *intelligent*—backed up by reasons and better situated with respect to the environment—rules for organizations.

Current attention for privacy matters offers an interesting test-bench for the previous considerations, as it requires to model norms about information use *in* niches and flows *between* niches (c.f. *contextual integrity*, (Nissenbaum, 2009)).

Even more recent is the call for *algorithm assurance* to ensure that privacy, fairness, explainability and contestability interests of stakeholders are served in automated decision making using machine learning techniques. Requirements depend on material conditions in the social niche. Skin color for instance matters when buying makeup. Gender matters to many in matching partners for dates. But for hiring decisions both are usually taboo. Explainability and fairness can be implemented through AI techniques, but the appropriateness of solutions must be weighed

against privacy considerations as they may indirectly expose sensitive information (e.g. (Chang and Shokri, 2020)).

Addressing privacy and algorithm assurance requires modeling norms, systematizing attention to failures, and adopting a diagnosis point of view towards organizations and organizational goals.

In the light of the conceptualizations presented in the previous sections, a starting point for this innovation is a sufficiently rich way of describing and reasoning about social roles, as for instance captured by Hohfeldian and CAES positions. Actually, because normative positions enter within the reasoning cycle as expectations, CAES descriptions offer a complete frame within which we expect role players to be “boundedly rational”, and for this reason they can be used to characterize the problems—of modeling, design, planning, monitoring, diagnosis, assessment, see e.g. (Breuker, 1994)—that the agents must solve in a certain social role, and the knowledge resources they require for that.

The harder the problem (from a C-D gap perspective), the harder is to verify as an observer that the agent is doing as expected, and for this reason the agent will need to argue its choices to other participants in the environment to explain that he acted correctly. Evidently, the agent *in that role* is biased in its information collection by its CAES description requirements, and therefore, to correctly check *ex-post* whether he took good decisions, we may need to rely for rationality at a higher level on the dialectical process taking place between the agents. A typical solution, compatible with Heiner’s theory of predictable behaviour, would be to come up with an increasingly rigorous *burden of proof* protocol for justifying decisions. If, on the other hand, performance can be easily scored (e.g. on recognizing faces, for instance) the depth of the reasoning process to be articulated can be minimal. This brings us to the design problem of deciding whether to utilize *tacit knowledge* methods (e.g. based on statistics) or *explicit knowledge* methods. Evidently, the first are for their nature (typically feed-forward) faster than the second. But there is something more than performance at stake.

For Watson-like systems, the results of the question-answering “race” amongst competing agents are clear: the judge-user is supposed to know the answer, there will be no requests about “why” that answer is correct. On similar lines, consider a mobile phone unlocking application based on face recognition. On the other hand, when a diagnostic device settles on the conclusion “the patient has the appendicitis”, it is natural for the user to demand *why*, expecting reasons that are acceptable to him. Similarly, a public administration cannot (or at least should not) deny e.g. a parking license or citizenship without explaining *why*. The *social* role demanded of the intelligent agents in the two types of applications is different, and this difference is at the root of the recent calls for *explainable AI* (Core et al., 2006). When AI devices have social relevance, they cannot neglect to provide reasons (determined by their role) for their functioning.

To conclude, if we solve the social-role acquisition problem, we can solve the explainable AI problem, because we will be able to identify higher-order burden of proof (and possibly protocols) that can be used to distribute computation to possibly specialized agents. The explainable AI problem is at the root of our problems

in designing complex, adaptive AI systems, as illustrated by the confused use of knowledge structures in the history of AI. If we solve the social-role acquisition problem, we can also improve the functioning of our institutions, because we would have a better environmental model in which to implement and test new institutional mechanisms, and in which to interpret social data.

Automating the social role acquisition problem from scratch would require the acquiring agent to be embedded in the social system in the same—or a very similar—way humans are. This is at present not a realizable condition. A first step towards solving the social-role acquisition problem is however realizable: considering collective requirements as those that communities reify in institutions, and applying the lessons that can be learned there in engineering.

## References

- Alchourrón, C. E. and E. Bulygin (1971). *Normative systems*. Springer-Verlag.
- Austin, J. L. (1975). *How to Do Things with Words*. Harvard University Press.
- Bardone, E (2011). *Seeking Chances: From Biased Rationality to Distributed Cognition*. Springer Berlin Heidelberg.
- Bench-Capon, T. J. M. et al. (2012). “A history of AI and Law in 50 papers: 25 years of the international conference on AI and Law”. In: *Artificial Intelligence and Law 20.3*, pp. 215–319.
- Boer, A. (2009). *Legal Theory, Sources of Law and the Semantic Web*. IOS Press.
- (2014). “Punishments, rewards, and the production of evidence”. In: *Proceedings of the 26th international conference on Legal Knowledge and Information Systems (JURIX 2014)*. Vol. FAIA 271. IOS Press, pp. 97–102.
- Boer, A. and T. van Engers (2011a). “An Agent-based Legal Knowledge Acquisition Methodology for Agile Public Administration”. In: *ICAIL 2011: The Thirteenth International Conference on Artificial Intelligence and Law*.
- (2011b). “Diagnosis of Multi-Agent Systems and Its Application to Public Administration”. In: *Business Information Systems Workshop*. Vol. LNBIP 97, pp. 258–269.
- Bostrom, N. (2003). “Ethical Issues in Advanced Artificial Intelligence”. In: *Cognitive, Emotive and Ethical Aspects of Decision Making in Humans and in Artificial Intelligence 2*, pp. 12–17.
- Bratman, M (1987). *Intention, plans, and practical reason*. Harvard University Press.
- Breuker, J. (1994). “Components of problem solving and types of problems”. In: *Proceedings of the International Conference on Knowledge Engineering and Knowledge Management (EKAW1994)* LNCS 867, pp. 118–136.
- Brooks, R. A. (1991). “Intelligence without representation”. In: *Artificial Intelligence 47.1987*, pp. 139–159.

- Chaitin, G. (2005). “Epistemology as Information Theory: From Leibniz to ”. In: *Alan Turing Lecture on Computing and Philosophy, E-CAP’05, European Computing and Philosophy Conference*.
- Chang, H. and R. Shokri (2020). “On the Privacy Risks of Algorithmic Fairness”. In: *eprint arXiv:2011.03731*. arXiv: 2011.03731.
- Core, M., H. Lane, M. Van Lent, D. Gomboc, S. Solomon, and M. Rosenberg (2006). “Building explainable artificial intelligence systems”. In: *Proceedings of the international conference on Innovative Applications of Artificial Intelligence (IAAI) 21.2*, pp. 1766–1773.
- Dari-Mattiacci, G. and G. D. Geest (2013). “The Rise of Carrots and the Decline of Sticks”. In: *University of Chicago Law Review* 80.1, pp. 341–392.
- De Filippi, P., C. Wray, and G. Sileno (2020). “Smart Contracts”. In: *Glossary of Distributed Technologies*. Ed. by V. Ferrari, A. Leiter, M. Mannan, and M.-C. Valiente. Internet Policy Review.
- De Landa, M. (2006). *A new philosophy of society: Assemblage theory and social complexity*. Continuum Intl Pub Group.
- De Mulder, R. and A. Meijer (2012). “From Trias to Tetras Politica: The Need for Additional Checks and Balances. Illustrated by Immigration Law”. In: *Public Administration in the Information Age: Revisited* 19, pp. 36–45.
- Dean, J. and S. Ghemawat (2008). “MapReduce: Simplified Data Processing on Large Clusters”. In: *Communication of the ACM* 51.1, pp. 107–113.
- Deleuze, G. and F. Guattari (1980). *Mille plateaux*. Éditions de Minuit.
- Dennett, D. C. (1987). *The Intentional Stance*. Cambridge, Massachusetts: MIT Press.
- Dessalles, J. L. (2013). “Algorithmic simplicity and relevance”. In: *Algorithmic probability and friends* 7070 LNAI, pp. 119–130.
- Dowe, D. L., J. Hernández-Orallo, and P. K. Das (2011). “Compression and intelligence: Social environments and communication”. In: *Proceedings of the 8th International Conference on Artificial General Intelligence*. Vol. 6830 LNAI. LNAI 9205, pp. 204–211.
- Ferraiolo, D., J. Barkley, and R. Kuhn (1992). “Role-Based Access Controls”. In: *ACM Transactions on Information and System Security* 2.1, pp. 34–64.
- Ferrucci, D. et al. (2010). “Building Watson: An Overview of the DeepQA Project”. In: *AI Magazine* 31.3, p. 59.
- Fonagy, P. and M. Target (1997). “Attachment and reflective function: their role in self-organization”. In: *Development and psychopathology* 9.4, pp. 679–700.
- Goodfellow, I. J., J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio (2014). “Generative Adversarial Nets”. In: *Advances in Neural Information Processing Systems (NIPS2014)*.
- Heiner, R. (1983). “The origin of predictable behavior”. In: *The American economic review* 73.4, pp. 560–595.
- Hohfeld, W. N. (1913). “Some Fundamental Legal Conceptions as Applied in Judicial Reasoning”. In: *Yale Law Journal* 23.1, pp. 16–59.
- Hutter, M. (2006). *The Human knowledge compression prize*. URL: <http://prize.hutter1.net>.

- Kahneman, D. and A. Tversky (1979). “Prospect Theory: An Analysis of Decision under Risk”. In: *Econometrica* 47, pp. 263–291.
- Legg, S. and M. Hutter (2007). “Universal intelligence: A definition of machine intelligence”. In: *Minds and Machines* 17.4, pp. 391–444. ISSN: 09246495. arXiv: 0712.3329.
- Lenat, D. B. (1995). “CYC: A Large-scale Investment in Knowledge Infrastructure”. In: *Communication of the ACM* 38.11, pp. 33–38.
- Mercier, H. and D. Sperber (2011). “Why do humans reason? Arguments for an argumentative theory”. In: *Behavioral and brain sciences* 34.2, pp. 57–74.
- Minsky, M. (1988). *Society of mind*. Simon & Schuster.
- Morgenstern, L. (1996). “The problem with solutions to the frame problem”. In: *The Robot’s Dilemma Revisited: The Frame Problem in AI*. Norwood, NJ: Ablex Publishing Corp., pp. 1–40.
- Newell, A. (1982). “The Knowledge Level”. In: *Artificial Intelligence* 18.1, pp. 87–127.
- Newell, A., J. C. Shaw, and H. A. Simon (1959). “Report on a general problem-solving program”. In: *Proceedings of the International Conference on Information Processing*, pp. 256–264.
- Newell, A. and H. A. Simon (1976). “Computer Science As Empirical Inquiry: Symbols and Search”. In: *Communication of the ACM* 19.3, pp. 113–126.
- Nissenbaum, H. (2009). *Privacy In Context: Technology Policy And The Integrity Of Social Life*. Stanford University Press, p. 304.
- Philippopoulos-Mihalopoulos, A. (2014). “Critical Autopoiesis and the Materiality of Law”. In: *International Journal for the Semiotics of Law* 27.2, pp. 389–418. ISSN: 15728722.
- Pollock, J. L. (1975). *Knowledge and Justification*. Princeton University Press.
- Pottage, A. (2012). “The Materiality of What?” In: *Journal of Law and Society* 39.1, pp. 167–183.
- Rao, A. S. and M. P. Georgeff (1995). “BDI Agents: From Theory to Practice”. In: *Proceeding of the 1st International Conference on Multi-agent Systems (ICMAS)*, pp. 312–319.
- Sartor, G. (2006). “Fundamental legal concepts: A formal and teleological characterisation”. In: *Artificial Intelligence and Law* 14.1, pp. 101–142.
- Seddon, J. (2008). *Systems Thinking in the Public Sector*. Triarchy Press.
- Sileno, G. (2016). “Aligning Law and Action”. PhD thesis. University of Amsterdam.
- Sileno, G., A. Boer, and T. van Engers (2014). “On the Interactional Meaning of Fundamental Legal Concepts”. In: *Proceedings of the 27th International Conference on Legal Knowledge and Information Systems (JURIX 2014)*. Vol. FAIA 271, pp. 39–48.
- (2015). “Commitments, Expectations, Affordances and Susceptibilities: Towards Positional Agent Programming”. In: *Proceedings fo the 18th Conferences on Principles and Practice of Multi-Agent Systems (PRIMA 2015)*. Vol. LNCS 9387, pp. 687–696.

- Sileno, G., A. Boer, and T. van Engers (2017). “Reading Agendas Between the Lines, an Exercise”. In: *Artificial Intelligence & Law* 25.1.
- (2018). “The Role of Normware in Trustworthy and Explainable AI”. In: *Proceedings of the XAILA workshop on eXplainable AI and Law, in conjunction with JURIX 2018*. arXiv: 1812.02471.
- (2020). “Monitoring and enforcement as a second-order guidance problem”. In: *JURIX 2020: The 33rd Annual Conference on Legal Knowledge and Information Systems*.
- Sileno, G., A. Saillenfest, and J.-L. Dessalles (2017). “A Computational Model of Moral and Legal Responsibility via Simplicity Theory”. In: *Proceedings of the 30th Int. Conf. on Legal Knowledge and Information Systems (JURIX 2017)* FAIA 302, pp. 171–176.
- Sileno, G., A. Boer, G. Gordon, and B. Rieder (2020). “Like Circles in the Water: Responsibility as a System-Level Function”. In: *Proceedings of the 3rd XAILA workshop on eXplainable AI and Law, in conjunction with JURIX 2020*.
- Silver, D. et al. (2016). “Mastering the Game of Go without Human Knowledge”. In: *Nature* 529.7587, pp. 484–489.
- Sowa, J. F. (1995). “Top-level Ontological Categories”. In: *International Journal Human-Computational Studies* 43.5-6, pp. 669–685.
- Sperber, D and D Wilson (1986). *Relevance: Communication and Cognition*. Wiley.
- Tin Kam Ho (1995). “Random decision forests”. In: *Proceedings of 3rd International Conference on Document Analysis and Recognition*. Vol. 1, 278–282 vol.1.
- Twining, W. (2010). “Normative and Legal Pluralism: A Global Perspective”. In: *Duke Journal of Comparative & International Law* 20.2000, pp. 473–517.
- Walton, D. N. (1988). “Burden of proof”. In: *Argumentation* 2.2, pp. 233–254.