

Multi-Agent Systems

Long Story Short

Andrea Omicini

<mailto:andrea.omicini@unibo.it>

Dipartimento di Informatica – Scienza e Ingegneria (DISI)

ALMA MATER STUDIORUM – Università di Bologna

AlxIA 2026

Perugia, Italy, 7 October 2026



- 1 Setting the Scene
- 2 Before the Name: Distributed AI
- 3 The Theory Years
- 4 The Field
- 5 The Quiet Decade
- 6 Agentic AI
- 7 The Next Decade



Next in Line...

- 1 Setting the Scene
- 2 Before the Name: Distributed AI
- 3 The Theory Years
- 4 The Field
- 5 The Quiet Decade
- 6 Agentic AI
- 7 The Next Decade



A Call for Papers

Modelling Autonomous Agents in a Multi-Agent World

- nothing in that title would look out of place in the programme of any agentic-AI venue of today
 - *autonomy*, *agents*, multiplicity, a *world* to *act* in
- ? so, when was it written?



1989

Demazeau & Müller (eds.), “Decentralized A.I.”^[Demazeau and Müller, 1990]

Proceedings of the First European Workshop on Modelling Autonomous Agents in a Multi-Agent World, Cambridge, England, 16–18 August 1989

- the *first* European workshop of the field
- the *name* of the field, at the time, was **distributed artificial intelligence (DAI)**
- ! thirty-seven years ago, the research question was already *written in full*

The Claim of this Tutorial

The question

- for more than thirty-five years the agent community has been answering *one* question
 - what is the right abstraction to encapsulate autonomy and intelligence in a distributed computational system?*
- and it produced the answer long *before* the technology could actually carry it
- *agentic AI* is the moment the technology arrived
 - *needing* exactly the concepts the community had already built
 - yet *rebuilding* several of them with limited understanding of the multi-agent system (MAS) literature

What this Tutorial is, and What is Not

- it is *not* a survey
 - no chronological march: chronology is the *substrate*; storylines are the *figure*
- it is a **claim with evidence**
 - every historical beat enters not because it happened, but because apparently it is *needed again now*



Agency = Autonomy $\times$ Intelligence

“AI as agency without intelligence”^[Floridi, 2023]

LLMs can process text with extraordinary success and often in a way that is indistinguishable from human output; however, they lack any intelligence, understanding, or cognitive ability

- the decoupling is *not* this speaker's invention
 - which is precisely what makes it safe to build on
- the mainstream has just discovered that the two axes *come apart*
 - something the agent community has known since it first had to define “agent” at all



A Coordinate System for Thirty-Five Years

- **agency** is the plane, with two independent axes
 - *autonomy*
 - who is in control, and of what
 - *intelligence*
 - what competence is brought to bear
- **encapsulation** is *not* a third axis
 - it is the *engineering discipline* that makes agency constructible
 - it is the reason the abstraction is called an “agent”, and not a “process”, an “object”, or a “model”

every era of this story can be plotted as a *region* on that plane

Plotting the Eras I

distributed artificial intelligence (DAI), 1989–1996

- autonomy asserted, under-specified; intelligence assumed, symbolic and social
- the agent is *being invented* as a unit of decomposition

agent theories, 1994–2001

- autonomy *formalised*; intelligence debated, symbolic against reactive
- agent = encapsulated control

Plotting the Eras II

AOSE and A&A, 2002–2010

- both axes stable; encapsulation *extended*
- environment, coordination, organisation get abstractions of their own

the deep-learning decade, 2010–2022

- autonomy dropped; intelligence maximised
- encapsulation dissolved
 - the *model* is the system

Plotting the Eras III

agentic AI, 2022–

- autonomy re-asserted, unspecified again; intelligence outsourced to the LLM
- encapsulation *ad hoc*: frameworks, not abstractions

the next decade

the *open problem*—and the subject of the *closing agenda*

Three Notions, Accumulated not Declared

- this tutorial does *not* open with a block of definitions
 - unlike many others I have given over the decades
- each notion is introduced when the *history forces it*, then re-used
 - **encapsulation**—Act I, and again in Act III
 - **autonomy**—Act II, and again in Act IV
 - **intelligence**—Act IV, and again in Act V
- ! the accumulation *is* the structure of the argument



The Plan

I *before the name*

- DAI, 1980s–1996

II *the theory years*

- 1994–2001

III *the field*

- 2002–2010

IV *the quiet decade*

- 2010–2022

V *agentic AI*

- 2022–

coda *the next decade*



Next in Line...

- 1 Setting the Scene
- 2 **Before the Name: Distributed AI**
- 3 The Theory Years
- 4 The Field
- 5 The Quiet Decade
- 6 Agentic AI
- 7 The Next Decade



The Problem Statement was Already Right

- everything the present moment calls new was in the *problem statement* of distributed artificial intelligence (DAI)
 - that intelligence is *distributed*
 - that it is *social*
 - that it is *situated*
 - that a single reasoner is the *wrong unit of account*
- ! the field did not lack the idea
 - it lacked the machine




MAAMAW, 1989–2001

- Modelling Autonomous Agents in a Multi-Agent World (MAAMAW), ten editions
 - 1989 Cambridge, 1990, 1991—published by Elsevier / North-Holland as *Decentralized A.I.* 1, 2, 3
[Demazeau and Müller, 1990] [Demazeau and Müller, 1991] [Werner and Demazeau, 1992]
 - from 1992 the Springer LNAI run: 830, 957, 1069, 1038, 1237, 1647
[Castelfranchi and Werner, 1994] [Castelfranchi and Müller, 1995] [Perram and Müller, 1996]
[Van de Velde and Perram, 1996] [Boman and Van de Velde, 1997] [Garijo and Boman, 1999]
 - 1992 is *Artificial Social Systems*, edited by Castelfranchi and Werner
- the European half of a field that also had North American DAI workshops and the Asian MAAC workshops

A Field that Merged Twice

- 1995—International Conference on Multi-Agent Systems (ICMAS), San Francisco, over 500 participants
 - itself the merger of the *European* MAAMAW workshops, the *Asian* MAAC workshops, and the *North American* distributed artificial intelligence (DAI) workshops
- four editions—San Francisco 1995, Kyoto 1996, Paris 1998, Boston 2000
- ! the field was already *organising itself* out of three continents, *before* it had a conference
- ! and it would merge again in 2002
 - see Act III



The Canon Arrives

- Bond & Gasser (eds.), *Readings in distributed artificial intelligence*, Morgan Kaufmann, 1988^[Bond and Gasser, 1988]
 - O'Hare & Jennings (eds.), *Foundations of distributed artificial intelligence*, Wiley, April 1996^[O'Hare and Jennings, 1996]
 - in the “Sixth-Generation Computer Technology” series
 - which is itself a period detail worth noticing
 - Weiss (ed.), *Multiagent Systems: A Modern Approach to distributed artificial intelligence*, MIT Press, 1999^[Weiss, 1999]
- ! a field with a **canon** is a field that expects *to last*



What a DAI Agent Is

O'Hare & Jennings, 1996 [O'Hare and Jennings, 1996]

...has an explicit representation of the world ... is situated within its environment ... solves a problem that requires intelligence ... deliberates / plans its course of actions ... is flexible ... is adaptable ... learns

? read that list again, and say which decade it comes from



Then and Now: the Feature List

Then—1996

world model, situatedness, problem solving requiring intelligence, deliberation and planning, flexibility, adaptability, learning

Now—2026

context and memory, environment and sandbox, reasoning, planning, tool use and re-planning, robustness, continual improvement

! item for item, a 1996 research programme and a 2026 product datasheet

Why Software Engineering Enters Here

- DAI posed an AI question
- but the answer it converged on – the *agent* – is first and foremost a *software-engineering* answer
- ! that observation is what this whole tutorial rests on

Decomposition & modularisation ^[Parnas, 1972]

*when software systems grow too complex, **decomposition** makes it possible to articulate them as ensembles of distinct, interacting components*

Encapsulation, Four Ways I

as a process^[Parnas, 1972]

the act of enclosing one or more items within a (physical or logical) container

in programming languages

a property of programming languages providing for mechanisms and abstractions allowing systems to be modularised around distinct components with limited interaction

Encapsulation, Four Ways II

in object-oriented programming

*some mechanisms to **rule access** to an object's **components**—plus some constructs promoting bundling of data with those methods operating on them*

in software engineering^[IEEE Standard Board, 1990]

*a software development technique that consists of **isolating** a system function or a set of data and operations on those data within a **module** [and] providing precise specifications for the module*

The Question

*What are the suitable software abstractions to **encapsulate** the many techniques available for the many diverse issues requiring artificial intelligence?*

What is the right place for AI techniques in intelligent systems engineering?

- this is Parnas's question, asked **one level up**
 - ! and it is still open in 2026
 - which is the whole point



The Answer

Agents

- agents are computational abstractions *encapsulating everything*
 - **computational autonomy**, for sure
 - and *intelligence*, in its many diverse forms
 - almost all scientific and technical communities have largely converged on that notion
- caveat** *not* on what an agent is—not on what *defines* an agent

! agents work as the **basic bricks** for the engineering of intelligent systems

Notion I: Encapsulation

Act I closes here

- the agent is *not* introduced as a metaphor for a person
 - it is introduced as a **unit of decomposition for AI**
 - the question is not “what is an agent like?” but “what does an agent *enclose*?”
- ! we will come back to encapsulation in Act III, with a harder question
- not *what does an agent encapsulate*, but *what else needs encapsulating*


Next in Line...

- 1 Setting the Scene
- 2 Before the Name: Distributed AI
- 3 The Theory Years**
- 4 The Field
- 5 The Quiet Decade
- 6 Agentic AI
- 7 The Next Decade



Seven Years of Defining Terms

- Agent Theories, Architectures, and Languages (ATAL), 1994–2001
 - eight editions, eight LNAI volumes
 - 890, [Wooldridge and Jennings, 1995b] 1037, [Wooldridge et al., 1996] 1193, [Müller et al., 1997] 1365, [Singh et al., 1998] 1555, [Müller et al., 1999] 1757, [Jennings and Lespérance, 2000] 1986, [Castelfranchi and Lespérance, 2001] 2333 [Meyer and Tambe, 2002]
 - it started as an *ECAI-94 workshop* in Amsterdam, 8–9 August 1994, edited by Wooldridge and Jennings, 24 papers
 - ! eight book spines in a row: a community building a theory *on purpose*
- in seven years the field did the conceptual work that the current wave is skipping
 - it fully elaborated its conceptual pillars
 - and consequently *defined its terms*



The Weak Notion of Agent

Weak agents are... [Wooldridge and Jennings, 1995a]

- *autonomous*
- *pro-active*
- *reactive*
 - to change
- *social*

- ! every “agentic” definition written since 2023 is a **weak-notion** definition
- usually with the same four properties, in a different order, with no citation

The Strong Notion of Agent

Strong agents... [Wooldridge and Jennings, 1995a]

- have *mental components*
 - belief, desire, intention, knowledge, ...
- the Belief, Desire, Intention (BDI) line follows: practical reasoning, means-ends reasoning, commitment strategies [Rao and Georgeff, 1995]
- AgentSpeak(L) makes it a programming language [Rao, 1996, Bordini et al., 2006]
- ! today's planner / executor / critic loops are **strong-notion architectures**
 - mostly without the hard part: *when* to (autonomously) *reconsider* an *intention*

Then and Now: Definitions

Then—1995

the *weak* and *strong* notions of agent

- behavioural properties on one side
- mental attitudes on the other
- a clean line between them

Now—2024—

competing definitions of “agentic”

- trade literature against MAS literature
- with the two senses used interchangeably
- often in the same paragraph

The Many Agents Around I

AI [Russell and Norvig, 2022]

the agent encapsulates *individual intelligence*—cognitive first, active second, despite *agens / agere* meaning “the one who acts”

distributed artificial intelligence [O'Hare and Jennings, 1996]

the agent encapsulates the *social dimension*—the individual *in a society*, distributed within a distributed environment



The Many Agents Around II

programming languages^[Odell, 2002]

- “an agent encapsulates *control flow*”
- “an agent is never *invoked*—it merely follows / drives its own control flow”

parallel & distributed systems^[Fuggetta et al., 1998]

mobility, and a reversal of perspective: “it is not that agents are mobile, it is that *objects are not*”

The Many Agents Around III

robotics^[Brooks, 1986]

situatedness—“a robot is a *physical* agent”; and intelligent behaviour with no symbolic representation at all

software engineering

the agent as an abstraction encapsulating information and knowledge, control, goal and task, intelligence, mobility—the basis of agent-oriented software engineering (AOSE)



The Many Agents Around IV

MAS^[Agre, 1995]

autonomy *and* interactivity: limited perception, limited capability, dependence on others and on the environment

each community reached for its own “agent” notion
to encapsulate *its own* hardest thing

Convergence on the Notion, not on the Definition

- AI, distributed artificial intelligence (DAI), programming languages, parallel & distributed systems, robotics, software engineering
 - six different roads, one destination
- they converge *on the notion...*


caveat ...not on the definition

- ! which is basically the position agentic AI is in today
 - and it is a *solvable* position, because it was solved once



The Counter-Current

- not everyone agreed that agency needed representation
 - Brooks
 - the subsumption architecture, intelligence *without* representation [Brooks, 1986, Brooks, 1991]
 - Suchman
 - *situated action*: plans are resources for action, not programs of it [Suchman, 1987]
 - Agre
 - computation and human experience, against the planning view [Agre, 1995]
- ! keep this thread in hand
 - in Act V, “representation is the bottleneck” will land as a *return*, not as a novelty



The Encapsulation Ladder [Odell, 2002]

- *monolithic* programming
 - nothing is encapsulated
 - *modular* programming
 - unit **behaviour** is
 - *object-oriented* programming
 - behaviour *and* **state**
 - *agent-oriented* programming
 - behaviour, state *and* **invocation**
- agents can say “Go!”
 - agents can also say “**No!**”

Data Flow Through Agents, Control Does Not

- in object-oriented programming, a message carries *data along with control*
- in agent-oriented programming, a message carries *data only*
- ! which is what lets information flow be designed *independently* of control flow

the boundary of an agent is a boundary that *control cannot cross*

Computational Autonomy

Agent

Agents are autonomous computational entities.

genus: agents are computational entities; differentia: agents are autonomous, in that they encapsulate control along with a criterion to govern it

Computational autonomy

agents are autonomous as they encapsulate (the thread of) control—control does not pass through agent boundaries; only data – knowledge, information – crosses agent boundaries

Lex Parsimoniae: Autonomy

Autonomy is the *only* defining feature

- agents have *no interface*
 - agents *cannot be controlled*
 - agents *cannot be invoked*
-
- everything else *follows*
 - *pro-activity*
 - “making something happen, rather than waiting for something to happen”; autonomous agents are pro-active *by definition*
 - *sociality, situatedness, reactivity, interactivity*
- ! multi-agent systems (MASs) are the only **legitimate containers** of autonomous agents



Autonomy Borrowed from Everywhere Else

- *etymology*

- Greek *autonomia* / *autonomos*, “having its own laws”; self-government, self-regulation, self-determination

- *philosophy, after Kant*

- “a more demanding notion of autonomy requires not only self-determination, but also **independence from oneself**”
- the ability to reason on and change one’s own *goals*, not just one’s actions

- *the military*

- automatic systems are pre-programmed and “not able to define the path according to some given goal”; autonomous ones are “self-directed toward a goal... governed by laws and strategies that direct their behaviour”

- *cars*

- six levels, 0 to 5, “from automatic to autonomous, one step at a time”^[SAE International, 2021]

Autonomy is Relational

Castelfranchi, 1995^[Castelfranchi, 1995]

autonomy has no meaning for an agent in isolation

- *goal-autonomous agent*
 - “an agent endowed with its own goals”
- *fully socially autonomous agent*^[Conte and Castelfranchi, 1995]
 - endogenous goals; decides among conflicting goals; adopts outside goals only through belief-mediated choice; self-interested
 - “it is impossible to change automatically the beliefs of an agent”
 - which is what protects its *cognitive autonomy*

Executive *vs.* Motivational Autonomy

Castelfranchi, 1995^[Castelfranchi, 1995]

- *executive autonomy*
 - given a goal, the agent is autonomous in achieving it by itself
 - *motivational autonomy*
 - the agent's goals are somehow *self-generated*, not externally imposed
- ! executive autonomy is not enough for real autonomy
- this is the most under-used weapon in the whole story
 - and it is about to become very useful indeed

The Test

Three questions to ask of anything called an agent

- does it *encapsulate control*?
 - does it have *no interface*?
 - is it *not invocable*?
-
- an LLM agent reached through a function-call API *is invoked*
 - its control passes through its own boundary, on every turn
 - ! by the *computational* definition it is *not* autonomous
 - in Castelfranchi's terms, it is *executively* autonomous at best, and *motivationally null*
 - ! that is a **precise critique**, not a complaint
 - and it is available only because the definitions were done thirty years ago

The Community Builds its Machinery

- *Autonomous Agents* (ACM), 1997–2001
 - Marina del Rey and after
- *PRIMA*, from 1998
 - the Pacific Rim
- *AgentLink*, from November 1998
 - “The European Network of Excellence for Agent-Based Computing”
- *European Agent Systems Summer School* (EASSS), from 1999
 - still running
- *COORDINATION*, from 1996
 - models and languages for (agent) interaction [Ciancarini et al., 1999]
- *WOA*, from 2000 [Mascardi and Omicini, 2026]
 - *Dagli Oggetti agli Agenti*, and it is still running too



Notion II: Autonomy

Act II closes here

- autonomy calls for *encapsulated control*, *plus a criterion to govern it*
 - it is the *only* defining feature
 - everything else is derived
 - it is *relational*: it has no meaning for an agent in isolation
 - it is *graded*, and *executive is not motivational*
- ! this is the conceptual *centre of gravity* of this tutorial
- everything in Act V is measured against it



Next in Line...

- 1 Setting the Scene
- 2 Before the Name: Distributed AI
- 3 The Theory Years
- 4 The Field**
- 5 The Quiet Decade
- 6 Agentic AI
- 7 The Next Decade




AAMAS 2002, Bologna

- three series end, one begins
 - the 6th *Autonomous Agents*
 - the 5th ICMAS
 - the 9th ATAL
- ! they become the 1st Autonomous Agents and Multi-Agent Systems (AAMAS)
 - general chairs: Gini and Ishida
 - programme chairs: Castelfranchi and Johnson
 - organisation chairs and treasurers: Omicini and Zambonelli
- ! we can adopt this event as the *conventional date* when MASs stopped being a *topic* and became a *discipline*



A Community that Never Stopped



Reading the Lifelines

- three foundational lines *stop* at 2001, one *starts* in 2002
- around it, an entire **workshop ecology** appears and persists
 - AOSE 2000–2012, ESAW 2000–2009, PROMAS and DALT 2003–2012
 - E4MAS 2004–2014, COIN 2006–2017, ESOA 2003–2006, SASO 2007–2019
 - EUMAS from 2003, WOA from 2000, EMAS from 2013, EXTRAAMAS from 2019
- ! and AIXIA runs across the whole picture, from 1988
 - the venue we are speaking at is *as old as the story*



AgentLink III

- an FP6 Coordination Action, 2004–2005, coordinated by Peter McBurney, Liverpool
- Work Package 5, the **Technical Forum**, was mine
- three forums
 - *TF1*—Rome, 30 June – 2 July 2004, over 100 researchers
 - *TF2*—Ljubljana, April 2005
 - *TF3*—Budapest, September 2005
- outputs as special issues in *KER* 20, ^[Omicini and McBurney, 2005] *Informatica* 29 ^[Omicini et al., 2005] and 30, ^[Omicini et al., 2006a] *TAAS* ^[Petta et al., 2007] and *JAAMAS* ^[Omicini and McBurney, 2008]



What the Field Thought Mattered, in 2004

- the nine thematic groups of Technical Forum I
 - AOSE, [Bernon et al., 2005a] PROMAS, [Dastani and Gómez-Sanz, 2005]
SELFORG, [Di Marzo Serugendo et al., 2005] TRUST, LEA, NETAGENTS,
HEALTHCARE, IIA4WE, BIOAGENTS [Luck and Merelli, 2005]

at least *six of those nine* are live agentic-AI concerns today

- methodology, programming, self-organisation, trust, learning, agents on the network
- ! a 2004 snapshot that reads as a perfectly-believable 2026 workshop call

The Field Becomes Teachable

- Weiss (ed.), MIT Press, 1999^[Weiss, 1999]
- Wooldridge, Wiley, 2002^[Wooldridge, 2002]
 - 2nd edition 2009^[Wooldridge, 2009]
- Vlassis, Morgan & Claypool, 2007^[Vlassis, 2007]
- Weyns, Springer, 2010^[Weyns, 2010]
- Dastani, Hindriks & Meyer (eds.), Springer, 2010^[Dastani et al., 2010]
- Ossowski (ed.), *Agreement Technologies*, Springer, 2013^[Ossowski, 2013]
- ! textbooks, a summer school, a conference, a journal, a European network
 - that is what a *discipline* looks like



The “Environment” Turn

- the field then solved a problem that agentic AI has not yet noticed it has

the environment is *not a backdrop*—it is something you *design*

- the environment as a **first-class abstraction** in MASs^[Weyns et al., 2007]
 - with its own abstractions, its own lifecycle, its own engineering^[Weyns and Holvoet, 2005]
- E4MAS, 2004–2014, is a workshop series about *exactly this*



Coordination Artefacts, 2004^[Omicini et al., 2004, Ricci et al., 2005]

Coordination artefacts


are artefacts used in the context of collaborative activities, mediating the interaction among actors involved in the same social context

- the lineage is explicit and borrowed on purpose
 - Activity Theory
 - every human activity is *mediated by mediating artefacts*^[Vygotskii, 1978]
 - Distributed Cognition
 - cognition is *distributed across individuals, tools and artefacts*^[Kirsh, 1999]
 - *cognitive artefacts*^[Norman, 1992]
 - “those artificial devices that maintain, display, or operate upon information in order to serve a representational function and that affect human cognitive performance”

Agens Faber, 2006 [Omicini et al., 2006b]

Toward a theory of artefacts for MASs

- the ability to *design and use tools* is a fundamental sign of intelligence [Hewes, 1993]
 - *homo faber* ties tool use to world-building, and to autonomy [Arendt, 1958]
 - ? so, is there the need for an *agens faber*?
-
- artefacts become a way of *programming* a MAS, [Ricci et al., 2006] and a way of *engineering its environment* [Viroli et al., 2005, Ricci and Viroli, 2005]
 - CArtaGO follows as the infrastructure [Ricci et al., 2009]
 - and the meta-model is consolidated in JAAMAS in 2008 [Omicini et al., 2008]



The A&A Meta-Model [Omicini et al., 2008] I

Agent

An A&A agent is an *autonomous computational entity*—*genus*: agents are computational entities; *differentia*: agents are autonomous, in that they encapsulate control along with a criterion to govern it

Artefact

An A&A artefact is a *computational entity aimed at the use by agents*—*genus*: artefacts are computational entities; *differentia*: artefacts are aimed to be used by agents

The A&A Meta-Model [Omicini et al., 2008] II

Cognitional artefact

An A&A cognitional artefact is an artefact aimed at the cognitive use by agents

MAS

An A&A MAS is a computational system made of agents and artefacts—genus: MAS is a computational system; differentia: its basic components are agents and artefacts



Artefacts Are Not Agents

- artefacts are *transparent, predictable, reactive, situated, non-autonomous*
- artefacts have *no goals*
 - they have a *function*, and, when used, a *destination*
 - a *destination* is “an external goal attached to an artefact by an agent, *in the act of using it*” [Conte and Castelfranchi, 1995]
- artefacts *do* have an interface
 - which is precisely what agents do not have

agents are designed to *govern*; artefacts are designed to *serve*

A Specification of Tool Use, 2006–2008^[Omicini et al., 2008]

Usage interface

the set of operations provided by an artefact

Operating instructions

*a description of the procedure an agent has to follow to meaningfully interact with an artefact **over time***

Function description

*what to **obtain** from an artefact—as distinct from operating instructions, which describe **how to exploit it***

Methodologies

- AOSE produced *methodologies*^[Bernon et al., 2005b]
 - with phases, abstractions, traceability
 - Gaia,^[Zambonelli et al., 2003] SODA,^[Omicini, 2001] Tropos^[Bresciani et al., 2004]
- SODA extended the methodology to the *environment*^[Molesini et al., 2006]
 - a roadmap to artefacts: design the workspace, not only the society
- organisations get their own abstractions too,^[Hübner et al., 2007] and a platform that puts the three together^[Boissier et al., 2013]
- ! remember this slide
 - the *absence* of its 2026 counterpart is the largest open gap in the coda



Notion I Again: What *Else* Needs Encapsulating

Act III closes here

everything that is *not autonomous*

- coordination
 - a medium that can be *programmed* [Omicini and Zambonelli, 1999]
- organisation
 - roles, groups, responsibilities, power
- norms and institutions
 - what may be done, and by whom
- the environment itself
 - workspaces, artefacts, tools
- ! the agent abstraction alone was never going to be enough
 - and the field *knew* that, and said so, twenty years ago



Next in Line...

- 1 Setting the Scene
- 2 Before the Name: Distributed AI
- 3 The Theory Years
- 4 The Field
- 5 The Quiet Decade**
- 6 Agentic AI
- 7 The Next Decade



Two Asymmetric Decades

- while mainstream AI went all-in on *learning*, [LeCun et al., 2015, Vaswani et al., 2017] the agent community kept the abstractions alive and kept *engineering* them
 - this act is short, and it is honest about the asymmetry of attention
- caveat the abstraction that the deep-learning decade offered for an intelligent system was “*the model*”
- no autonomy, no environment, no society, no encapsulation
 - one function, end to end



Consolidation

- PROMAS + AOSE + DALT converge into Engineering Multi-Agent Systems (EMAS), from 2013
 - three workshops become one, exactly as three conferences had in 2002
- AGERE!, 2011–2021, keeps the actors-and-agents conversation going with the programming-languages community
- ! a field that consolidates rather than expands is a field under less attention
 - not a field with less to say




Where the Energy Went

- *self-organisation and pervasive systems* [Di Marzo Serugendo et al., 2006]
 - SAPERE, and a decade of self-organising coordination [Zambonelli et al., 2015, Omicini and Viroli, 2011]
 - *stigmergy*
 - coordination through the medium, named in 1959 for termites [Grassé, 1959, Parunak, 2006]
- *MAS-based simulation*
 - as a *method*, with validation obligations [Herller and Klügl, 2006]
- *norms, institutions, socio-technical systems*—COIN, 2006–2017
- ! emergence is *engineered through the environment*, not hoped for from population size



XAI: Intelligence Back as a Design Question

- explainable artificial intelligence (XAI) ^[Rosenfeld and Richardson, 2019, Barredo Arrieta et al., 2020]
- EXTRAAMAS from 2019, xAI 2023–2024, and the EXPECTATION line
- symbolic knowledge *extraction* and *injection* ^[Ciatto et al., 2024]
 - a principled alternative to reading a trace and calling it an explanation
- ! this is where *intelligence* re-enters as a *design* question rather than a performance metric
 - not “how well does it score” but “what about it can be *inspected*”



Autonomy in Living Systems [Maturana and Varela, 1980]

Autopoiesis

the ability of a complex system to maintain its own overall coherence, in terms of structure and organisation, through the mutual interactions of its components

Autonomous systems

*they acquire the property of specifying **their own rules of behaviour**... they do not work as transducers or functions for converting input instructions into output products... they are the **sources of their own activity***

Heteronomous vs. Autonomous^[Thompson, 2010]

Heteronomous

*one whose organisation is defined by input–output information flow and **external** mechanisms of control*

Autonomous

*defined by its **endogenous**, self–organising, and self–controlling dynamics... it has input and output—which, **alone**, do not determine the system*

- ! a Turing machine is a **heteronomous** system; an agent is not
 - which is the deepest available statement of why “agent” is not a synonym for “program”

Where does Autonomy come from?

Autonomy of the Living^[Rosslenbroich, 2014]

Living systems are autonomous in the sense that they maintain themselves in form and function within time, and achieve a self-determined flexibility

- *increasing autonomy* is “an evolutionary shift in the system–environment relationship”
 - *interactive* autonomy
 - direct environmental influences gradually reduced
 - *constitutive* autonomy
 - self-referential intrinsic functions generated and enhanced
- ! *autonomy is not autarchy*
 - it is graded, and it is relative



Intelligence in Living Systems

- intelligence is a *long-term success story*
 - roughly four billion years of it
- *in the beginning it was not the word*
 - prokaryotes with no mind, no reasoning, and no representation nonetheless *sensed* others and their environment
 - “sensing is the most elementary sort of cognition”
- **non-explicit intelligence**
 - no logic, no reasoning, no representation
 - just an implicit competence curating life by the rules of *homeostasis*, “relentlessly executed autonomously according to no explicit directions”



Explicit and Non-Explicit

Damasio, 2021^[Damasio, 2021]

*Mind is one way of referring to the active production and display of **images** arising from actual perception, from memory recall, or from both... Images of every kind... are natural vehicles of knowledge*

- **quorum sensing**
 - coordination with neither mind nor representation^[Diggle et al., 2007]
 - a biological precedent for stigmergy, found in bacteria
- ! there is much more intelligence than just **rational** intelligence

Notion III: Intelligence

Act IV closes here

- intelligence is *not* the same as performance
 - it comes in *explicit* and *non-explicit* forms, and both work
 - it is a property of *living systems* that computing is still learning to formulate
- ! autonomy and intelligence are not software properties that happened to be invented in 1995
- ? and, how does intelligence connect with autonomy—in general, and within agents?

Next in Line...

- 1 Setting the Scene
- 2 Before the Name: Distributed AI
- 3 The Theory Years
- 4 The Field
- 5 The Quiet Decade
- 6 Agentic AI**
- 7 The Next Decade



The Technology Arrives

- agentic AI is the field's *vindication* and its *embarrassment* at once
 - *vindication*, because agents are now the central abstraction of AI, exactly as expected
 - *embarrassment*, because the re-invention is happening *without the theory*
- ! and part of the reason is a blind spot of our own, which this act will concede



Arriving at Agents by Deduction from a Limitation

Augmented language models^[Mialon et al., 2023]

the fundamental problem of LLMs is reliance on a single parametric model and a limited context, typically the n previous or surrounding tokens

the solution is integrating LMs with both reasoning and tools

- so “agents and action come into play”
- ! the mainstream arrives at agents by *deduction from a limitation*
 - which is a lovely historical irony, and worth naming out loud

From Stateless Prediction to Stateful Control

- a plain LLM is reducible to *sequence continuation*
 - $P(\text{next token} \mid \text{previous tokens})$
- an *agentic system* wraps that model inside a larger, *stateful*, *goal-directed*, *action-taking* process
 - ReAct,^[Yao et al., 2022] Toolformer,^[Schick et al., 2023] AutoGPT,^[Richards, 2023]
planner–executor, memory–augmented^[Lewis et al., 2020]
- they all share: iterative control, explicit state, action semantics, environmental feedback, adaptation over time
- ! “from text as *output*, to text as *control signal*”



The Terminological Turf War

- on one side, the agentic–AI trade literature^[Gullf, 2025, Ranjan et al., 2025, Biswas and Talukdar, 2025]
 - an *AI agent* is “a system designed to perceive its environment and take actions to achieve a specific goal... an evolution from a standard LLM, enhanced with the abilities to plan, use tools, and interact with its surroundings”
- on the other, the MAS engineering literature^[Collier et al., 2026]
 - ! a different interpretation for the same terms is apparently available, *and working*
 - ? what is the right one? *is* there a right one?



Do Not Resolve it by Fiat

- resolve it by *applying the test*, and showing what passes
 - encapsulated control
 - no interface
 - not invocable
 - e.g.
 - an agent behind a function-call schema fails all three
 - an agent with its own control loop, its own goals, and no entry point passes all three
 - ! this is not gatekeeping
 - it is the difference between a *definition* and a *marketing category*
- caveat and we need to understand each and every available and useful notion of agent and MAS, not only our own



Then and Now: Tools

Then—2004–2008

coordination artefacts; *agens faber*; the A&A triad of *usage interface*, *operating instructions*, *function description*

Now—2023—^[Qin et al., 2024]

LLM tool calling; tool schemas and registries; MCP-style tool exposure


- ! a tool needs a *function description*
 - what to obtain
- distinct from *operating instructions*
 - how to exploit it over time
- ! today's schemas **conflate the two**



Tools Have Not Been Invented by Agentic AI

one of the two accounts of tool use is a *theory*

- the A&A account says what an artefact *is*, what it is *for*, and how an agent relates to it
 - use, selection, construction, manipulation
 - and a *destination* attached in the act of using
- the schema account says what arguments a function takes
- ! both are useful: yet, only one supports composition, and provides a basis for *design*



Then and Now: Memory

Then—1992–2007^[Norman, 1992, Kirsh, 1999]

distributed cognition; cognitive artefacts; cognitive stigmergy

Now—2022—^[Lewis et al., 2020]

RAG; external and episodic memory; vector stores; memory governance

- ! cognition distributed across artefacts changes the *unit of analysis*
 - so memory design is a *social-system* question, not a retrieval optimisation



Then and Now: Orchestration and Guardrails

Then—2000–2017^[Hübner et al., 2007]

organisations, roles and groups; institutions; normative MASs; norms enforced *by the environment*

Now—2023—

multi-agent orchestration; supervisor patterns; guardrails, policy layers, permission scopes

- ! organisation is an abstraction with its own *semantics* – roles, responsibilities, power – rather than a call graph between prompts
- ! and *regimentation* is not *regulation*: a norm the environment enforces is not a norm the agent is asked to respect

Then and Now: Methodology

Then—2000–2012^[Zambonelli et al., 2003, Omicini, 2001, Bresciani et al., 2004]

AOSE: Gaia, SODA, Tropos—phases, abstractions, traceability, and a workshop series to argue about them in

Now—2024—^[Gullí, 2025]

framework recipes, prompt patterns, “agentic design patterns”


- ! there *were* methodologies
 - and no equivalent exists for agentic AI
- ! this is the **largest live gap**, and it is squarely in AIxIA’s competence

The Counter-Beat: Our Own Failure

FIPA-era standardisation was *logocentric*

it invested in *communication* – agent communication languages, speech acts, interaction protocols – and comparatively little in *pragmatic action*: what an agent does to the world, and with what [FIPA ACL, 2002]

- the A&A line was in part a *reaction* to that
 - but A&A did not become the standard; *Agent Communication Language* (*ACL*) did
- ! so when agentic AI needed a theory of *tool use*, the standard it inherited from the agent world had *almost nothing to say*
 - and tool calling was built from scratch



The Empirical Brake

- planning is *not* just *fluent plan-shaped text*
 - PlanBench: by obfuscating the domain vocabulary while preserving the logical structure of the plan, performance disappears ^[Valmeekam et al., 2023]
 - ARC-AGI: a stress test for *representation formation*, not a verdict on general intelligence
- the counter-current of Act II returns here, from the other side
 - Brooks said intelligence needs *less* representation; the benchmarks suggest agency needs *more explicit* representation
- ! *representation is the bottleneck*
- ! do not confuse *plan fluency* with *autonomous planning competence*

The Field's Own Current Answer

Mascardi & Omicini (eds.), 2026^[Mascardi and Omicini, 2026]

The Agents Journey: Twenty-Five Years of Multi-agent Systems

Springer, LNCS State-of-the-Art Survey 16395, April 2026—open access,
funded by ENGES

- multi-agent systems (MASs) are “at the core of the emerging Agentic AI field, providing the conceptual and computational foundations for building intelligent systems composed of interacting autonomous entities”
- two chapters are the forward vectors of this tutorial
 - ch. 12, from symbolic to neurosymbolic (NeSy) systems^[Agiollo et al., 2026]
 - ch. 14, AI infrastructure, from gigastructure to edge intelligence^[Omicini et al., 2026]



Where the Three Notions Stand

Act V closes here

- agentic AI *has* intelligence
 - borrowed
- it *asserts* autonomy
 - unspecified
- it *lacks* encapsulation
 - frameworks instead of abstractions

! that sentence is the whole diagnosis

! and each of its three clauses is a research agenda



Next in Line...

- 1 Setting the Scene
- 2 Before the Name: Distributed AI
- 3 The Theory Years
- 4 The Field
- 5 The Quiet Decade
- 6 Agentic AI
- 7 The Next Decade**




Not “We Were Right”

We always expected MASs to become central.

So the interesting question is no longer whether we were right—it is which of our answers still hold, and which we got wrong.

- so, we should look for a *research programme*, not a victory lap
- my draft here: *one agenda per notion*, each stated as a problem you could work on



Agenda I: Encapsulation

A unified theory of agents and tools

- there is *no* unified agent–tool theory
- the A&A account covers agents using artefacts; the agentic–AI account covers models calling functions
- ？ what covers *both*, and lets them compose?

A methodology

- the 2000s produced Gaia, SODA, Tropos
- the 2020s have produced frameworks and prompt patterns
- ？ what is the agent–oriented software engineering (AOSE) of agentic AI?

Agenda II: Autonomy

Levels of autonomy for agentic systems

- agentic AI discusses autonomy *informally*
- other fields have graded taxonomies
 - six levels for driving automation, [SAE International, 2021] four for military systems
- and the agent community has the distinction that neither of those has: **executive vs. motivational** [Castelfranchi, 1995]
- an LLM agent given a goal and left to achieve it is *executively* autonomous and *motivationally null*
- ! a levels-of-autonomy scale built on encapsulated control and the executive / motivational axis is concrete, publishable, and obviously useful
- ! it is the single most **actionable** thing this tutorial can hand you

Agenda III: Intelligence

Neurosymbolic integration and intermediate representations

- the position: from symbolic to neurosymbolic (NeSy) systems, as a quest for *integration* [Agiollo et al., 2026]
 - the argument: the benchmarks, and the bottleneck [Valmeekam et al., 2023]
 - the slogan: **representation before autonomy**
-
- ? which representations must be *explicit* for an agent to be *governable*?
 - ? and can those representations be *induced*, rather than hand-designed?
 - intermediate representations are not implementation details
 - they are what makes bias *contestable* and agency *governable* [Ciatto et al., 2025]

Crossing All Three: Infrastructure

From gigastructure to edge intelligence^[Omicini et al., 2026]

- *where agents run* is becoming a first-class question again
- which is Act III's *environment turn* returning
 - at the scale of the deployment substrate, rather than the workspace
- ! the environment was a *workspace* in 2004, and it is a *continent-scale infrastructure* in 2026
- ! and it is still **the same abstraction problem**




Long Story Short

- the question has not changed since 1989
 - *what is the right abstraction to encapsulate autonomy and intelligence in a distributed computational system?*
- **autonomy** was defined, and the definition still discriminates
- **intelligence** turned out to be separable from agency
 - and plural
- **encapsulation** is the part still being rebuilt
 - quite badly, right now

Multi-agent systems paved the way

The road is still the same road

Read the WOA Version



- these slides

talks.apice.unibo.it/MasAixia2026

The Agents Journey: Twenty-Five Years of Multi-agent Systems

- Springer, LNCS State-of-the-Art Survey 16395, April 2026 [Mascardi and Omicini, 2026]
- *open access*

<https://link.springer.com/10.1007/978-3-032-22940-3>

Multi-Agent Systems

Long Story Short

Andrea Omicini

<mailto:andrea.omicini@unibo.it>

Dipartimento di Informatica – Scienza e Ingegneria (DISI)

ALMA MATER STUDIORUM – Università di Bologna

AlxIA 2026

Perugia, Italy, 7 October 2026



References I

[Agiollo et al., 2026] Agiollo, A., Calegari, R., Ciatto, G., Magnini, M., Omicini, A., and Sabbatini, F. (2026)

Intelligent agents from symbolic to neurosymbolic systems: The quest for integration

In [Mascardi and Omicini, 2026], chapter 12, pages 320–339

APICe DOI:10.1007/978-3-032-22940-3_12

[Agre, 1995] Agre, P. E. (1995)

Computational research on interaction and agency

Artificial Intelligence, 72(1–2):1–52

Special volume on computational research on interaction and agency, part 1

APICe DOI:10.1016/0004-3702(94)00054-5

[Arendt, 1958] Arendt, H. (1958)

The Human Condition

University of Chicago Press, Chicago, IL, USA

APICe

[Barredo Arrieta et al., 2020] Barredo Arrieta, A., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., and Herrera, F. (2020)

Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI

Information Fusion, 58:82–115

APICe DOI:10.1016/j.inffus.2019.12.012

[Bernon et al., 2005a] Bernon, C., Cossentino, M., and Pavón, J. (2005a)

Agent-oriented software engineering

In [Omicini and McBurney, 2005], pages 99–116

APICe DOI:10.1017/S0269888905000421

[Bernon et al., 2005b] Bernon, C., Cossentino, M., and Pavón, J. (2005b)

An overview of current trends in European AOSE research

In [Omicini et al., 2005], pages 379–390

APICe DOI:10.31449/inf.v29i4.54



References II

- [Biswas and Talukdar, 2025] Biswas, A. and Talukdar, W. (2025)
Building Agentic AI Systems
 Packt Publishing
<https://www.oreilly.com/library/view/building-agentic-ai/9781803238753/>
- [Boissier et al., 2013] Boissier, O., Bordini, R. H., Hübner, J., Ricci, A., and Santi, A. (2013)
Multi-agent oriented programming with JaCaMo
Science of Computer Programming, 78(6):747–761
 Special section on Agent-oriented Design Methods and Programming Techniques for Distributed Computing in Dynamic and Complex Environments
 APICe DOI:10.1016/j.scico.2011.10.004
- [Boman and Van de Velde, 1997] Boman, M. and Van de Velde, W., editors (1997)
Multi-Agent Rationality, volume 1237 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg
 APICe DOI:10.1007/3-540-63077-5
- [Bond and Gasser, 1988] Bond, A. H. and Gasser, L., editors (1988)
Readings in Distributed Artificial Intelligence
 Morgan Kaufmann, San Mateo, CA, USA
 APICe
- [Bordini et al., 2006] Bordini, R. H., Braubach, L., Dastani, M., El Fallah Seghrouchni, A., Gómez-Sanz, J. J., Leite, J., O'Hare, G., Pokahr, A., and Ricci, A. (2006)
A survey of programming languages and platforms for multi-agent systems
 In [Omicini et al., 2006a], pages 33–44
 APICe <https://www.informatica.si/index.php/informatica/article/view/71>



References III

- [Bresciani et al., 2004] Bresciani, P., Giorgini, P., Giunchiglia, F., Mylopoulos, J., and Perini, A. (2004)
Tropos: An agent-oriented software development methodology
Autonomous Agents and Multi-Agent Systems, 8(3):203–236
 APiCe DOI:10.1023/B:AGNT.0000018806.20944.ef
- [Brooks, 1986] Brooks, R. A. (1986)
Achieving artificial intelligence through building robots
 A.I. Memo 899, Massachusetts Institute of Technology (MIT), Cambridge, MA, USA
 APiCe <http://hdl.handle.net/1721.1/6451>
- [Brooks, 1991] Brooks, R. A. (1991)
Intelligence without reason
 In Mylopoulos, J. and Reiter, R., editors, *12th International Joint Conference on Artificial Intelligence (IJCAI 1991)*, volume 1, pages 569–595, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
 APiCe <http://dl.acm.org/citation.cfm?id=1631258>
- [Castelfranchi, 1995] Castelfranchi, C. (1995)
Guarantees for autonomy in cognitive agent architecture
 In [Wooldridge and Jennings, 1995b], pages 56–70
 APiCe DOI:10.1007/3-540-58855-8_3
- [Castelfranchi and Lespérance, 2001] Castelfranchi, C. and Lespérance, Y., editors (2001)
Intelligent Agents VII, volume 1986 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg
 APiCe DOI:10.1007/3-540-44631-1
- [Castelfranchi and Müller, 1995] Castelfranchi, C. and Müller, J.-P., editors (1995)
From Reaction to Cognition, volume 957 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg
 APiCe DOI:10.1007/BFb0027051



References IV

- [Castelfranchi and Werner, 1994] Castelfranchi, C. and Werner, E., editors (1994)
Artificial Social Systems, volume 830 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg
 APICe DOI:10.1007/3-540-58266-5
- [Ciancarini et al., 1999] Ciancarini, P., Omicini, A., and Zambonelli, F. (1999)
Coordination technologies for Internet agents
Nordic Journal of Computing, 6(3):215–240
 APICe
- [Ciatto et al., 2025] Ciatto, G., Aguzzi, G., Battistini, R., Baiardi, M., Burattini, S., and Ricci, A. (2025)
Exploiting genai for plan generation in bdi agents
 In Lynce, I., Murano, N., Vallati, M., Villata, S., Chesani, F., Milano, M., Omicini, A., and Dastani, M., editors, *28th European Conference on Artificial Intelligence, 25–30 October 2025, Bologna, Italy – Including 14th Conference on Prestigious Applications of Intelligent Systems (PAIS 2025)*, volume 413 of *Frontiers in Artificial Intelligence and Applications*, pages 3495–3502, Amsterdam, The Netherlands. IOS Press
 APICe DOI:10.3233/FAIA251223
- [Ciatto et al., 2024] Ciatto, G., Sabbatini, F., Agiollo, A., Magnini, M., and Omicini, A. (2024)
Symbolic knowledge extraction and injection with sub-symbolic predictors: A systematic literature review
ACM Computing Surveys, 56(6):1–35
 APICe DOI:10.1145/3645103
- [Collier et al., 2026] Collier, R., Mascardi, V., and Ricci, A., editors (2026)
Agents and Multi-Agent Systems Development: Platforms, Toolkits, Technologies
 Springer
 APICe DOI:10.1007/978-3-032-01082-7



References V

- [Conte and Castelfranchi, 1995] Conte, R. and Castelfranchi, C. (1995)
Cognitive and Social Action
 Garland Science
 APiCe DOI:10.4324/9780203783221
- [Damasio, 2021] Damasio, A. R. (2021)
Feeling and Knowing: Making Minds Conscious
 Pantheon Books, New York, NY, USA, 1st edition
 APiCe <https://www.penguinrandomhouse.com/books/572664/feeling-and-knowing-by-antonio-damasio/>
- [Dastani and Gómez-Sanz, 2005] Dastani, M. and Gómez-Sanz, J. J. (2005)
Programming multi-agent systems
 In [Omicini and McBurney, 2005], pages 151–164
 APiCe DOI:10.1017/S0269888905000482
- [Dastani et al., 2010] Dastani, M., Hindriks, K. V., and Meyer, J.-J. C., editors (2010)
Specification and Verification of Multi-agent Systems
 Springer, New York, NY, USA
 Foreword by Wiebe van der Hoek
 APiCe DOI:10.1007/978-1-4419-6984-2
- [Demazeau and Müller, 1990] Demazeau, Y. and Müller, J.-P., editors (1990)
Decentralized A.I.
 North-Holland / Elsevier Science Publishers B.V., Amsterdam, The Netherlands
 1st European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW 1989), Cambridge, England,
 16–18 August 1989. Proceedings. Title page dated 1990
 APiCe



References VI

- [Demazeau and Müller, 1991] Demazeau, Y. and Müller, J.-P., editors (1991)
Decentralized A.I. 2, Amsterdam, The Netherlands. North-Holland
 Proceedings of the second European Workshop on Modelling Autonomous Agents in a Multi-Agent World, Saint-Quentin en Yvelines, France, August 13–16, 1990. Also issued in paperback as ISBN 978-0-444-56779-6.
 APiCe <https://shop.elsevier.com/books/decentralized-ai-2/demazeau/978-0-444-89051-1>
- [Di Marzo Serugendo et al., 2005] Di Marzo Serugendo, G., Gleizes, M.-P., and Karageorgos, A. (2005)
Self-organization in multi-agent systems
 In [Omicini and McBurney, 2005], pages 165–189
 APiCe DOI:10.1017/S0269888905000494
- [Di Marzo Serugendo et al., 2006] Di Marzo Serugendo, G., Gleizes, M.-P., and Karageorgos, A. (2006)
Self-organisation and emergence in MAS: An overview
 In [Omicini et al., 2006a], pages 45–54
 APiCe <https://www.informatica.si/index.php/informatica/article/view/72>
- [Diggle et al., 2007] Diggle, S. P., Griffin, A. S., Campbell, G. S., and West, S. A. (2007)
Cooperation and conflict in quorum-sensing bacterial populations
Nature, 450(7168):411–414
 APiCe DOI:10.1038/nature06279
- [FIPA ACL, 2002] FIPA ACL (2002)
Agent Communication Language Specifications
 Foundation for Intelligent Physical Agents (FIPA)
 The fipa.org domain no longer serves FIPA; cited from the Internet Archive snapshot of 4 January 2018
 APiCe <https://web.archive.org/web/20180104053953/http://www.fipa.org/repository/aclspecs.html>



References VII

[Floridi, 2023] Floridi, L. (2023)

AI as agency without intelligence: On ChatGPT, large language models, and other generative models

Philosophy and Technology, 2023

APiCe DOI:10.2139/ssrn.4358789

[Fuggetta et al., 1998] Fuggetta, A., Picco, G. P., and Vigna, G. (1998)

Understanding code mobility

IEEE Transactions on Software Engineering, 24(5):342–361

APiCe DOI:10.1109/32.685258

[Garijo and Boman, 1999] Garijo, F. J. and Boman, M., editors (1999)

Multi-Agent System Engineering, volume 1647 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg

APiCe DOI:10.1007/3-540-48437-X

[Grassé, 1959] Grassé, P.-P. (1959)

La reconstruction du nid et les coordinations interindividuelles chez *Bellicositermes natalensis* et *Cubitermes* sp. la théorie de la stigmergie: Essai d'interprétation du comportement des termites constructeurs

Insectes Sociaux, 6(1):41–80

APiCe DOI:10.1007/BF02223791

[Gullí, 2025] Gullí, A. (2025)

Agentic Design Patterns: A Hands-On Guide to Building Intelligent Systems

Springer

DOI:10.1007/978-3-032-01402-3

[Herrler and Klügl, 2006] Herrler, R. and Klügl, F. (2006)

Simulation

In Kirn, S., Herzog, O., Lockemann, P., and Spaniol, O., editors, *Multiagent Engineering: Theory and Applications in Enterprises*, chapter V–2, pages 575–596. Springer Berlin Heidelberg

APiCe DOI:10.1007/3-540-32062-8_27



References VIII

[Hewes, 1993] Hewes, G. W. (1993)

A history of speculation on the relation between tools and languages

In Gibson, K. R. and Ingold, T., editors, *Tools, Language and Cognition in Human Evolution*, pages 20–31. Cambridge University Press

APICe <https://www.cambridge.org/us/academic/subjects/life-sciences/neuroscience/tools-language-and-cognition-human-evolution>

[Hübner et al., 2007] Hübner, J. F., Sichman, J. S., and Boissier, O. (2007)

Developing organised multiagent systems using the MOISE+ model: programming issues at the system and agent levels

International Journal of Agent Oriented Software Engineering, 1(3/4):370–395

APICe DOI:10.1504/IJAOSE.2007.016266

[IEEE Standard Board, 1990] IEEE Standard Board (1990)

IEEE standard glossary of software engineering terminology

IEEE Std 610.12–1990, IEEE, New York, NY, USA

Reaffirmed 12–9–2002

APICe DOI:10.1109/IEEESTD.1990.101064

[Jennings and Lespérance, 2000] Jennings, N. R. and Lespérance, Y., editors (2000)

Intelligent Agents VI, volume 1757 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg

APICe DOI:10.1007/10719619

[Kirsh, 1999] Kirsh, D. (1999)

Distributed cognition, coordination and environment design

In Bagnara, S., editor, *3rd European Conference on Cognitive Science (ECCS'99)*, pages 1–11, Certosa di Pontignano, Siena, Italy.

Istituto di Psicologia, Consiglio Nazionale delle Ricerche

APICe <https://philarchive.org/archive/KIRDCC>



References IX

- [LeCun et al., 2015] LeCun, Y., Bengio, Y., and Hinton, G. (2015)
Deep learning
Nature, 521(7553):436–444
 APiCe DOI:10.1038/nature14539
- [Lewis et al., 2020] Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-t., Rocktäschel, T., Riedel, S., and Kiela, D. (2020)
Retrieval-Augmented Generation for knowledge-intensive NLP tasks
 In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H., editors, *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474. Curran Associates, Inc.
 APiCe <https://proceedings.neurips.cc/paper/2020/hash/6b493230205f780e1bc26945df7481e5-Abstract.html>
- [Luck and Merelli, 2005] Luck, M. and Merelli, E. (2005)
Agents in bioinformatics
 In [Omicini and McBurney, 2005], pages 117–125
 APiCe DOI:10.1017/S0269888905000433
- [Mascardi and Omicini, 2026] Mascardi, V. and Omicini, A., editors (2026)
The Agents Journey, volume 16395 of *Lecture Notes in Computer Science*
 Springer Nature Switzerland, Cham
 APiCe DOI:10.1007/978-3-032-22940-3
- [Maturana and Varela, 1980] Maturana, H. R. and Varela, F. G. (1980)
Autopoiesis and Cognition: The Realization of the Living, volume 42 of *Boston Studies in the Philosophy of Science*
 D. Reidel Publishing Company, Dordrecht, Holland
 APiCe DOI:10.1007/978-94-009-8947-4
- [Meyer and Tambe, 2002] Meyer, J.-J. C. and Tambe, M., editors (2002)
Intelligent Agents VIII, volume 2333 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg
 APiCe DOI:10.1007/3-540-45448-9



References X

- [Mialon et al., 2023] Mialon, G., Dessì, R., Lomeli, M., Nalmpantis, C., Pasunuru, R., Raileanu, R., Rozière, B., Schick, T., Dwivedi-Yu, J., Celikyilmaz, A., Grave, E., LeCun, Y., and Scialom, T. (2023)
Augmented language models: a survey
Transactions on Machine Learning Research
 APiCe <https://openreview.net/forum?id=jh7wH2AzKK>
- [Molesini et al., 2006] Molesini, A., Omicini, A., Denti, E., and Ricci, A. (2006)
SODA: A roadmap to artefacts
 In Dikenelli, O., Gleizes, M.-P., and Ricci, A., editors, *Engineering Societies in the Agents World VI*, volume 3963 of *LNAI*, pages 49–62. Springer
 6th International Workshop (ESAW 2005), Kuşadası, Aydın, Turkey, 26–28 October 2005. Revised Selected and Invited Papers
 APiCe DOI:10.1007/11759683_4
- [Müller et al., 1999] Müller, J. P., Singh, M. P., and Rao, A. S., editors (1999)
Intelligent Agents V, volume 1555 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg
 APiCe DOI:10.1007/3-540-49057-4
- [Müller et al., 1997] Müller, J. P., Wooldridge, M. J., and Jennings, N. R., editors (1997)
Intelligent Agents III, volume 1193 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg
 APiCe DOI:10.1007/BFb0013568
- [Norman, 1992] Norman, D. A. (1992)
Design principles for cognitive artifacts
Research in Engineering Design, 4(1):43–50
 APiCe DOI:10.1007/BF02032391
- [Odell, 2002] Odell, J. (2002)
Objects and agents compared
Journal of Object Technologies, 1(1):41–53
 APiCe http://www.jot.fm/issues/issue_2002_05/column4



References XI

- [O'Hare and Jennings, 1996] O'Hare, G. M. and Jennings, N. R., editors (1996)
Foundations of Distributed Artificial Intelligence
 Sixth-Generation Computer Technology. John Wiley & Sons Ltd., hardcover edition
 APiCe <https://www.wiley.com/en-gb/Foundations+of+Distributed+Artificial+Intelligence-p-9780471006756>
- [Omicini, 2001] Omicini, A. (2001)
SODA: Societies and infrastructures in the analysis and design of agent-based systems
 In Ciancarini, P. and Wooldridge, M. J., editors, *Agent-Oriented Software Engineering*, volume 1957 of *Lecture Notes in Computer Science*, pages 185–193. Springer-Verlag
 1st International Workshop (AOSE 2000), Limerick, Ireland, 10 June 2000. Revised Papers
 APiCe DOI:10.1007/3-540-44564-1_12
- [Omicini and McBurney, 2005] Omicini, A. and McBurney, P. (2005)
Special issue “Revised reports from selected Technical Forum Groups at the 1st & 2nd AgentLink III Technical Forums”
The Knowledge Engineering Review, 20(2):95–196
 APiCe <https://www.cambridge.org/core/journals/knowledge-engineering-review/issue/C66711DC522542960D4F5EE362BDA3C7>
- [Omicini and McBurney, 2008] Omicini, A. and McBurney, P. (2008)
Special issue on foundations, advanced topics and industrial perspectives of multi-agent systems
Autonomous Agents and Multi-Agent Systems, 17(3):367–475
 APiCe <https://link.springer.com/journal/10458/17/3>
- [Omicini et al., 2005] Omicini, A., Petta, P., and Gams, M. (2005)
Special issue “hot topics in european agent research I”
Informatica, 29(4):377–443
 APiCe <https://www.informatica.si/index.php/informatica/issue/view/117>




References XII

- [Omicini et al., 2006a] Omicini, A., Petta, P., and Gams, M. (2006a)
Special issue “hot topics in european agent research II”
Informatica, 30(1):1–82
 APiCe <https://www.informatica.si/index.php/informatica/issue/view/118>
- [Omicini et al., 2026] Omicini, A., Ricci, A., and Mascardi, V. (2026)
AI infrastructure: From gigastructure to edge intelligence with multi-agent systems
 In [Mascardi and Omicini, 2026], chapter 14, pages 366–384
 APiCe DOI:10.1007/978-3-032-22940-3_14
- [Omicini et al., 2006b] Omicini, A., Ricci, A., and Viroli, M. (2006b)
Agens Faber: Toward a theory of artefacts for MAS
Electronic Notes in Theoretical Computer Sciences, 150(3):21–36
 1st International Workshop “Coordination and Organization” (CoOrg 2005), COORDINATION 2005, Namur, Belgium, 22 April 2005. Proceedings
 APiCe DOI:10.1016/j.entcs.2006.03.003
- [Omicini et al., 2008] Omicini, A., Ricci, A., and Viroli, M. (2008)
Artifacts in the A&A meta-model for multi-agent systems
Autonomous Agents and Multi-Agent Systems, 17(3):432–456
 Special Issue on Foundations, Advanced Topics and Industrial Perspectives of Multi-Agent Systems
 APiCe DOI:10.1007/s10458-008-9053-x
- [Omicini et al., 2004] Omicini, A., Ricci, A., Viroli, M., Castelfranchi, C., and Tummolini, L. (2004)
Coordination artifacts: Environment-based coordination for intelligent agents
 In Jennings, N. R., Sierra, C., Sonenberg, L., and Tambe, M., editors, *3rd international Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004)*, volume 1, pages 286–293, New York, USA. ACM
 APiCe DOI:10.1109/AAMAS.2004.10070



References XIII

- [Omicini and Viroli, 2011] Omicini, A. and Viroli, M. (2011)
Coordination models and languages: From parallel computing to self-organisation
The Knowledge Engineering Review, 26(1):53–59
 Special Issue 01 (25th Anniversary Issue)
 APICe DOI:10.1017/S026988891000041X
- [Omicini and Zambonelli, 1999] Omicini, A. and Zambonelli, F. (1999)
Coordination for Internet application development
Autonomous Agents and Multi-Agent Systems, 2(3):251–269
 Special Issue: Coordination Mechanisms for Web Agents
 APICe DOI:10.1023/A:1010060322135
- [Ossowski, 2013] Ossowski, S., editor (2013)
Agreement Technologies, volume 8 of *Law, Governance and Technology Series*
 Springer, Dordrecht, The Netherlands
 APICe DOI:10.1007/978-94-007-5583-3
- [Parnas, 1972] Parnas, D. L. (1972)
On the criteria to be used in decomposing systems into modules
Communications of the ACM, 15(12):1053–1058
 APICe DOI:10.1145/361598.361623
- [Parunak, 2006] Parunak, H. V. D. (2006)
A survey of environments and mechanisms for human-human stigmergy
 In Weyns, D., Parunak, H. V. D., and Michel, F., editors, *Environments for Multi-Agent Systems II*, volume 3830 of LNCS, pages
 163–186. Springer
 APICe DOI:10.1007/11678809_10



References XIV

- [Perram and Müller, 1996] Perram, J. W. and Müller, J.-P., editors (1996)
Distributed Software Agents and Applications, volume 1069 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg
 APiCe DOI:10.1007/3-540-61157-6
- [Petta et al., 2007] Petta, P., Omicini, A., Payne, T. R., and McBurney, P. (2007)
Special issue on the AgentLink III Technical Forums
ACM Transactions on Autonomous and Adaptive Systems, 2(4)
 APiCe DOI:10.1145/1293731
- [Qin et al., 2024] Qin, Y., Liang, S., Ye, Y., Zhu, K., Yan, L., Lu, Y., Lin, Y., Cong, X., Tang, X., Qian, B., Zhao, S., Hong, L., Tian, R., Xie, R., Zhou, J., Gerstein, M., Li, D., Liu, Z., and Sun, M. (2024)
ToolLLM: Facilitating large language models to master 16000+ real-world APIs
In 12th International Conference on Learning Representations (ICLR 2024)
 APiCe <https://openreview.net/forum?id=dHng200Jjr>
- [Ranjan et al., 2025] Ranjan, S., Chembachere, D., and Lobo, L. (2025)
Agentic AI in Enterprise: Harnessing Agentic AI for Business Transformation
 Springer
 DOI:10.1007/979-8-8688-1542-3
- [Rao, 1996] Rao, A. S. (1996)
AgentSpeak(L): BDI agents speak out in a logical computable language
In Van de Velde, W. and Perram, J. W., editors, Agents Breaking Away, volume 1038 of *Lecture Notes in Computer Science*, pages 42–55. Springer, Berlin, Heidelberg
 7th European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW'96), Eindhoven, The Netherlands, 22–25 January 1996, Proceedings
 APiCe DOI:10.1007/BFb0031845

References XV

[Rao and Georgeff, 1995] Rao, A. S. and Georgeff, M. P. (1995)

BDI agents: From theory to practice

In Lesser, V. R. and Gasser, L., editors, *1st International Conference on Multi Agent Systems (ICMAS 1995)*, pages 312–319, San Francisco, CA, USA. The MIT Press

APiCe <https://www.aaai.org/Papers/ICMAS/1995/ICMAS95-042.pdf>

[Ricci et al., 2009] Ricci, A., Piunti, M., Viroli, M., and Omicini, A. (2009)

Environment programming in CArAgO

In Bordini, R. P., Dastani, M., Dix, J., and El Fallah Seghrouchni, A., editors, *Multi-Agent Programming II: Languages, Platforms and Applications*, Multiagent Systems, Artificial Societies, and Simulated Organizations, chapter 8, pages 259–288. Springer

APiCe DOI:10.1007/978-0-387-89299-3_8

[Ricci and Viroli, 2005] Ricci, A. and Viroli, M. (2005)

Coordination artifacts: A unifying abstraction for engineering environment-mediated coordination in MAS

In [Omicini et al., 2005], pages 433–443

APiCe <https://www.informatica.si/index.php/informatica/article/view/59>


[Ricci et al., 2005] Ricci, A., Viroli, M., and Omicini, A. (2005)

Environment-based coordination through coordination artifacts

In Weyns, D., Parunak, H. V. D., and Michel, F., editors, *Environments for Multi-Agent Systems*, volume 3374 of *LNAI*, pages 190–214. Springer

1st International Workshop (E4MAS 2004), New York, NY, USA, 19 July 2004. Revised Selected Papers

APiCe DOI:10.1007/b106134



References XVI

[Ricci et al., 2006] Ricci, A., Viroli, M., and Omicini, A. (2006)

Programming MAS with artifacts

In Bordini, R. P., Dastani, M., Dix, J., and El Fallah Seghrouchni, A., editors, *Programming Multi-Agent Systems*, volume 3862 of *LNAI*, pages 206–221. Springer

3rd International Workshop (PROMAS 2005), AAMAS 2005, Utrecht, The Netherlands, 26 July 2005. Revised and Invited Papers

APiCe DOI:10.1007/11678823_13

[Richards, 2023] Richards, T. B. (2023)

Auto-GPT: An autonomous GPT-4 experiment

GitHub Repository

<https://github.com/DataBassGit/Auto-GPT>

[Rosenfeld and Richardson, 2019] Rosenfeld, A. and Richardson, A. (2019)

Explainability in human-agent systems

Autonomous Agents and Multi-Agent Systems, 33(6):673–705

APiCe DOI:10.1007/s10458-019-09408-y

[Rosslenbroich, 2014] Rosslenbroich, B. (2014)

On the Origin of Autonomy. A New Look at the Major Transitions in Evolution

History, Philosophy and Theory of the Life Sciences. Springer International Publishing

APiCe DOI:10.1007/978-3-319-04141-4

[Russell and Norvig, 2022] Russell, S. J. and Norvig, P. (2022)

Artificial Intelligence: A Modern Approach

Pearson Series on Artificial Intelligence. Pearson Education Limited, global edition, 4th edition

APiCe [https://www.pearson.com/uk/educators/higher-education-educators/program/](https://www.pearson.com/uk/educators/higher-education-educators/program/Russell-Artificial-Intelligence-A-Modern-Approach-Global-Edition-4th-Edition/PGM100003084094.html)

Russell-Artificial-Intelligence-A-Modern-Approach-Global-Edition-4th-Edition/PGM100003084094.html



References XVII

- [SAE International, 2021] SAE International (2021)
Taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles
 Surface Vehicle Recommended Practice J3016C, SAE International
 Issued 2014-01-16. Revised 2016-09-30; 2018-06-15; 2021-04-30
 APICe https://www.sae.org/standards/j3016_202104-taxonomy-definitions-terms-related-driving-automation-systems-road-motor-vehicles/
- [Schick et al., 2023] Schick, T., Dwivedi-Yu, J., Dessì, R., Raileanu, R., Lomeli, M., Hambro, E., Zettlemoyer, L., and Lewis, M. (2023)
Toolformer: Language models can teach themselves to use tools
 In 37th Conference on Neural Information Processing Systems (NeurIPS 2023)
 APICe <https://openreview.net/forum?id=Yacmpz84TH>
- [Singh et al., 1998] Singh, M. P., Rao, A. S., and Wooldridge, M. J., editors (1998)
Intelligent Agents IV, volume 1365 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg
 APICe DOI:10.1007/BFb0026744
- [Suchman, 1987] Suchman, L. A. (1987)
Plans and Situated Actions: The Problem of Human-Machine Communication
 Cambridge University Press, New York, NYU, USA
 APICe
- [Thompson, 2010] Thompson, E. (2010)
Mind in Life: Biology, Phenomenology, and the Sciences of Mind
 Belknap Press
 APICe <http://www.hup.harvard.edu/catalog.php?isbn=9780674057517>



References XVIII

- [Valmeekam et al., 2023] Valmeekam, K., Marquez, M., Olmo, A., Sreedharan, S., and Kambhampati, S. (2023)
PlanBench: An extensible benchmark for evaluating large language models on planning and reasoning about change
 In *37th Conference on Neural Information Processing Systems (NeurIPS 2023) – Datasets and Benchmarks Track*
 APiCe https://proceedings.neurips.cc/paper_files/paper/2023/hash/7a92bcdede88c7afd108072faf5485c8-Abstract-Datasets_and_Benchmarks.html
- [Van de Velde and Perram, 1996] Van de Velde, W. and Perram, J. W., editors (1996)
Agents Breaking Away, volume 1038 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg
 APiCe DOI:10.1007/BFb0031841
- [Vaswani et al., 2017] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017)
Attention is all you need
 In Guyon, I., von Luxburg, U., Bengio, S., Wallach, H. M., Fergus, R., Vishwanathan, S., and Garnett, R., editors, *Advances in Neural Information Processing Systems (NIPS 2017)*, volume 30, pages 5998–6008. Curran Associates, Inc.
 APiCe <https://papers.nips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html>
- [Viroli et al., 2005] Viroli, M., Omicini, A., and Ricci, A. (2005)
Engineering MAS environment with artifacts
 In Weyns, D., Parunak, H. V. D., and Michel, F., editors, *2nd International Workshop “Environments for Multi-Agent Systems” (E4MAS 2005)*, pages 62–77, AAMAS 2005, Utrecht, The Netherlands
 APiCe
- [Vlassis, 2007] Vlassis, N. (2007)
A Concise Introduction to Multiagent Systems and Distributed Artificial Intelligence, volume 2 of *Synthesis Lectures on Artificial Intelligence and Machine Learning*
 Morgan & Claypool
 APiCe DOI:10.2200/S00091ED1V01Y200705AIM002

References XIX

- [Vygotskiĭ, 1978] Vygotskiĭ, L. S. (1978)
Mind in Society: Development of Higher Psychological Processes
 Harvard University Press, Cambridge, MA, USA
 APiCe <http://www.hup.harvard.edu/catalog/VYGMIX.html>
- [Weiss, 1999] Weiss, G., editor (1999)
Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence
 The MIT Press, Cambridge, MA, USA
 APiCe <https://mitpress.mit.edu/9780262731317/multiagent-systems/>
- [Werner and Demazeau, 1992] Werner, E. and Demazeau, Y., editors (1992)
Decentralized A.I. 3, Amsterdam, The Netherlands. North-Holland
 Proceedings of the Third European Workshop on Modelling Autonomous Agents in a Multi-Agent World, Kaiserslautern, Germany, August 5-7, 1991
 APiCe <https://www.amazon.com/dp/0444896619>
- [Weyns, 2010] Weyns, D. (2010)
Architecture-Based Design of Multi-Agent Systems
 Springer-Verlag Berlin Heidelberg, Berlin, Heidelberg
 Foreword by Len Bass
 APiCe DOI:10.1007/978-3-642-01064-4
- [Weyns and Holvoet, 2005] Weyns, D. and Holvoet, T. (2005)
On the role of environments in multiagent systems
 In [Omicini et al., 2005], pages 409-421
 APiCe <https://www.informatica.si/index.php/informatica/article/view/57>



References XX

- [Weyns et al., 2007] Weyns, D., Omicini, A., and Odell, J. (2007)
Environment as a first-class abstraction in multi-agent systems
Autonomous Agents and Multi-Agent Systems, 14(1):5–30
 Special Issue on Environments for Multi-agent Systems
 APiCe DOI:10.1007/s10458-006-0012-0
- [Wooldridge and Jennings, 1995a] Wooldridge, M. and Jennings, N. R. (1995a)
Intelligent agents: Theory and practice
Knowledge Engineering Review, 10(2):115–152
 APiCe DOI:10.1017/S0269888900008122
- [Wooldridge, 2002] Wooldridge, M. J. (2002)
An Introduction to MultiAgent Systems
 John Wiley & Sons Ltd., Chichester, UK, 1st edition
 APiCe
- [Wooldridge, 2009] Wooldridge, M. J. (2009)
An Introduction to MultiAgent Systems
 John Wiley & Sons Ltd., Chichester, UK, 2nd edition
 APiCe <http://www.wiley.com/en-it/An+Introduction+to+MultiAgent+Systems,+2nd+Edition-p-9780470519462>
- [Wooldridge and Jennings, 1995b] Wooldridge, M. J. and Jennings, N. R., editors (1995b)
Intelligent Agents, volume 890 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg
 APiCe DOI:10.1007/3-540-58855-8
- [Wooldridge et al., 1996] Wooldridge, M. J., Müller, J. P., and Tambe, M., editors (1996)
Intelligent Agents II, volume 1037 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg
 APiCe DOI:10.1007/3-540-60805-2



References XXI

[Yao et al., 2022] Yao, S., Zhao, J., Yu, D., Shafran, I., Narasimhan, K. R., and Cao, Y. (2022)

ReAct: Synergizing reasoning and acting in language models

In *NeurIPS 2022 Foundation Models for Decision Making Workshop*

APiCe <https://openreview.net/forum?id=tvI4u1ylcqs>

[Zambonelli et al., 2003] Zambonelli, F., Jennings, N. R., and Wooldridge, M. (2003)

Developing multiagent systems: The Gaia methodology

ACM Transactions on Software Engineering and Methodology (TOSEM), 12(3):317–370

APiCe DOI:<http://doi.acm.org/10.1145/958961.958963>

[Zambonelli et al., 2015] Zambonelli, F., Omicini, A., Anzenberger, B., Castelli, G., DeAngelis, F. L., Di Marzo Serugendo, G., Dobson, S., Fernandez-Marquez, J. L., Ferscha, A., Mamei, M., Mariani, S., Molesini, A., Montagna, S., Nieminen, J., Pianini, D., Risoldi, M., Rosi, A., Stevenson, G., Viroli, M., and Ye, J. (2015)

Developing pervasive multi-agent systems with nature-inspired coordination

Pervasive and Mobile Computing, 17:236–252

Special Issue “10 years of Pervasive Computing” In Honor of Chatschik Bisdikian

APiCe DOI:[10.1016/j.pmcj.2014.12.002](https://doi.org/10.1016/j.pmcj.2014.12.002)

