

# Does Agentic AI merely incarnate Semantic Web Services?

Anonymous ACL submission

## Abstract

The recent publication of agentic AI protocols and systems (such as MCP, A2A, etc.) has sparked tremendous excitement in the AI community, as it enables agents to use and combine tools and services in order to achieve complex goals. This very promise is not new though, and has been made over 20 years ago by an initiative called “Semantic Web Services”. In this paper, we point out that (1) the underlying ideas of agentic AI can already be found in the Semantic Web Services vision, and (2) that abstractions, protocols, data models, and descriptions developed in the Web and Semantic Web community within and outside the (Semantic) Web Services strand of research can be beneficially applied in today’s endeavour to realise agentic AI. At the same time, we argue that looking back to history reveals certain still unresolved challenges.

## 1 Introduction

Soon after the invention of the Web by Tim Berners-Lee in the early ’90s of the last century, the emergence of tools and services available online have sparked the idea of *automated agents* running and orchestrating such services on our behalf, seamlessly achieving users’ goals and supporting our daily lives. This idea is manifested in the renowned “Semantic Web” paper originally published in 2001 in the Scientific American and recently reprinted (Berners-Lee et al., 2023). The very idea behind this original Semantic Web vision was based on well-defined protocols and metadata standards for describing Web content and services in a declarative, interoperable fashion, enabling agents to align and compose services on the fly. The first proposals for such protocols were based on then state-of-the-art Web

services technologies (Alonso et al., 2004), suggesting to combine those with the respective Semantic Web metadata standards (Martin et al., 2005; Roman et al., 2005). Yet, despite helping to conceptualise the building blocks of interoperable, machine-processable service ecosystems, these initiatives, dubbed “Semantic Web services”, have hardly been adopted in practice.

With the recent advent of generative AI models, and their underlying transformer models (Vaswani et al., 2017), that do not only allow to generate language, but also code, solve mediation and integration problems, even can learn how to use external tools (Schick et al., 2023), and plan courses of action to solve complex goals (Yao et al., 2023), we see the ideas of agentic AI refuelled. Likewise, frameworks and protocols to make generative AI models interact with such tools and services have been put in operation and are gaining traction rapidly, such as Anthropic’s Model Context Protocol (MCP) (Anthropic, 2025) and Google’s Agent2Agent protocol (A2A) (Google and A2A Project Contributors, 2025). As such, this may be a good time to ask ourselves whether indeed all the roadblocks that have prevented earlier successes in agentic AI have been successfully overcome with these recent technological advances?

To this end, the present paper aims at putting the historic efforts towards Semantic Web Services and recent trends and advances in agentic AI side by side. Our goal is to highlight useful “historical” conceptualisations, discuss in how far these are paralleled or incarnated by current agentic AI frameworks, and finally spotlight still unresolved challenges, with potential routes ahead.

In the remainder of this paper we first re-

cap base terminology and concepts from the (Semantic) Web Services Literature, in order to scope what we mean by “service” and other components needed in a truly agentic AI and service ecosystem. Next, we discuss in more detail in how far these concepts are implemented by recent agentic AI protocols. We conclude with challenges in terms of what we consider (still) unresolved challenges.

## 2 Web Services Terminology

The Web Services (Alonso et al., 2004) and Semantic Web Services (Fensel et al., 2006; Martin et al., 2005; Kopecký et al., 2007) literature has introduced several core concepts for building agentic ecosystems. We briefly recall these concepts and attempt to scope them for our purposes, in order to reflect which of them have been implemented by both preceeding technologies and current “agentic AI” protocols.

**Services, Operations and Functions** We herein assume a *service* as a collection of *service operations* accessible via a common access point, or *endpoint*, operated by a remote *agent*, which provides access to the service operations via an *interface*, that expose the services’s *capabilities* to the outside. A simple common abstraction of service operations is a *remote procedure call*.

Assuming such a service representing some agent, when taking this responsible agent’s capabilities into account, one could also distinguish between fully observable (locally) executable operations (using the agent’s own computing power and operating on data from the local state only) as *functions*, as opposed to *remote operations*, that require external services. Petrie called such functions “subroutines”(Petrie, 2016, Section 1.1.5); where we note here that also functions/subroutines could be discoverable and advertised by a service: a recent public effort to provide such a library service for discoverable functions is Wikifunctions.<sup>1</sup> As opposed to functions, true remote (service) operations, require capabilities *beyond the agent’s local control or observability*, according to Preist’s definition (Preist, 2004) of a service as the provision of an (abstract) value or *product* in the context of the

domain of application through interactions, by one party to another. This lack of observability also typically necessitates some form of agreement or *contract*, before actual delegation to remote service operations, potentially achieved through *negotiation*, i. e., interactions preceding the value delivery, related to non-functional aspects (see below), such as the alignment of the part-taking agents’ policies (Kampik et al., 2022).

**Provider and Requester; Invocation** In the context of a service ecosystem, the literature typically distinguishes between *service provider*, as for the agent providing a service and respective operations to other agents, and *service requester*, as for an agent that *invokes* (or: requests) services, beyond its own capabilities or for other reasons preferably delegated. It is important to note that this distinction, while putting the requester’s goal in the focus, largely ignores the provider’s goal(s).

**Service Descriptions** Services can be described for different purposes, e. g. to facilitate discovery, composition, binding and execution, as outlined below. In the simplest case, service descriptions contain overall *capabilities* of the service, its exposed operations as well as syntactical *interface* descriptions of *inputs* and *outputs* of specific operations, next to the address of the endpoint, and necessary authentication. These more functional aspects may be augmented by a description of *non-functional aspects*. We may understand those aspects as – abstract or explicit - (not necessarily fully observable) *conditions* and *rules*, over the internal states of different agents relevant to the service ecosystems (requester, provider, legislator, etc.), for instance *preconditions* that need to be fulfilled to successfully invoke a particular service operation, or *postconditions* that are guaranteed upon such successful invocation. Some even argued to include behavioural information in the description, i. e. in which order different operations should be invoked. Non-functional aspects may additionally refer to e. g. policies and regulations to execute an operation (legally) compliant, payment modalities, service level agreements or alike. In sum, non-functional aspects refer to (not fully observable) *conditions* and *rules*, over the internal states of different agents (re-

<sup>1</sup><https://www.wikifunctions.org/>




Figure 1: The Publish-Find-Invoke Pattern

183 requester, provider) or also rules governing the  
 184 service ecosystems as a whole (legislator, or-  
 185 ganization etc.).

186 **Discovery** For services and their operations  
 187 to be invocable, their availability must be  
 188 known to potential requesters. The process  
 189 of acquiring knowledge about relevant services  
 190 is called *service discovery*. On the service  
 191 provider side, the facilitation of discovery may  
 192 involve advertisement, in the form of either  
 193 decentralised self-advertisement or *publishing*  
 194 service descriptions into dedicated *reposito-*  
 195 *ries* or *registries*, which represent separate ser-  
 196 vices. To *find* the right services' descriptions  
 197 for downstream tasks including composition  
 198 and invocation, this discovery requires *search*  
 199 capabilities of *service descriptions*. Related  
 200 problems include those of *service matchmak-*  
 201 *ing* (to determine which of the discovered ser-  
 202 vices fulfil the requester's requirements) and  
 203 *selection* (to finally choose from the discovered  
 204 services). Those steps can be facilitated by  
 205 methods including string similarity measures  
 206 or ontological reasoning.

207 Corresponding registries facilitate the im-  
 208 plementation of the publish-find-invoke (or: -  
 209 bind) pattern for composition (Endrei et al.,  
 210 2004; Papazoglou and van den Heuvel, 2007),  
 211 which is commonplace in many service-  
 212 oriented approaches, including UPnP, Web  
 213 Services, and CORBA (Aiello, 2018), see Fig-  
 214 ure 1. In this pattern a Provider publishes de-  
 215 scriptions about the services they provide at a  
 216 Registry, where a Requester can find descrip-  
 217 tions, and use the descriptions' information for  
 218 invocations.

219 **Composition, Choreography, Orchestra-**  
 220 **tion** In case a single service call does not  
 221 achieve a given goal, a composition of different  
 222 available services and service operations may  
 223 achieve the said goal.

224 Establishing a viable composition to goal  
 225 completion may involve/ be preceded by *ne-*  
 226 *gotiations* establishing compliance with the




Figure 2: Service Orchestration. An orchestrator makes two sequential calls (solid) to two services, and receives corresponding responses (dashed).




Figure 3: Service Choreography. Two services call (solid) and respond to (dashed) each other sequentially.

227 rules and conditions –yet in practice negoti-  
 228 ations hardly happen and customers are left  
 229 with standardised service-level agreements to  
 230 which they can agree or not–, and *planning*  
 231 the sequence of invocations in the correct or-  
 232 der. This planning may be done by humans  
 233 who manually craft an orchestration, or by  
 234 automated planners –yet this automation is  
 235 not done at scale, probably due to the high  
 236 requirements of such planners regarding for-  
 237 malisation of corresponding goals and seman-  
 238 tic web service descriptions.

239 A composition may be described in differ-  
 240 ent ways, in terms of an explicit *workflow* or  
 241 *process*, that decomposes the overall goal into  
 242 subtasks down to single service or function  
 243 calls –often called an *orchestration*, see Fig-  
 244 ure 2–, or more declarative means which rather  
 245 define a certain order or parallel execution of  
 246 services, or, more generally, a *protocol* –often  
 247 called *choreography*, see Figure 3.

248 **Execution, Binding** Finally, for practical  
 249 effect, a composition of services needs get *ex-*  
 250 *ecuted*. To this end, values, services, and com-  
 251 munication channels need to be chosen (also  
 252 called *binding*), and the services invoked, all  
 253 governed by the composition. Depending on  
 254 the nature of the composition, orchestration  
 255 or choreography, there is or is not necessar-  
 256 ily a central entity to govern the composition.  
 257 In case of the orchestration, this entity is the  
 258 orchestrating service requester, in case of a  
 259 choreography there may or may not be a ded-

260	icated middleware in place that can serve as	XSD – descriptions of the service input and	308
261	message broker or even check compliance to	output messages’ grammar.	309
262	the protocol.		
263	<b>Mediation</b> is the task of resolving mis-	<b>Discovery</b> UDDI (Universal Description	310
264	matches in data (e.g. terminology or mes-	Discovery and Integration) as an early stan-	311
265	sage format mismatches), processes (e.g. be-	dard to advertise services and implement a re-	312
266	havioural mismatches), or functionality (e.g.	spective <i>registry</i> , distinguished between differ-	313
267	capability mismatches) between different ser-	ent levels of description relating to general ser-	314
268	vices to enable them to work together. Me-	vice description, such as contact information	315
269	diation can occur during search, composition,	(White pages), business description (in terms	316
270	and invocation.	of e.g. USPSC business/sector categories),	317
271	In the following, we will reflect in how far	and technical descriptions, e.g., linking to a	318
272	traditional Web service technologies, as well as	WSDL web service description. Although pub-	319
273	Semantic Web services have implemented the	licly running UDDI registries were operated	320
274	envisioned automation of the above-mentioned	(e.g. by IBM, Microsoft and SAP), all of these	321
275	steps, followed by a discussion of agentic pro-	seem to have stopped operation before 2010.	322
276	ocols.		
277	<b>3 Web Services Technologies</b>	<b>Orchestration</b> The Web Services Business	323
278	We introduce core technologies from the WS-*	Process Execution Language (Alves et al.,	324
279	stack, and Semantic Web Services, and briefly	2007) is an XML-based standard that allows	325
280	summarise.	to describe an orchestration of services in a	326
281	<b>3.1 WS-* Technologies</b>	process, i.e. calls to service operations, made	327
282	The Web services stack proposed to enable	by one single service requester, put into order	328
283	automation mostly through technologies like	by some control flow (e.g. sequential, parallel	329
284	SOAP, WSDL, and UDDI, providing stan-	ordering). For software that executes service	330
285	dardized protocols, interface descriptions and	orchestrations there are many commercial of-	331
286	registries (with limited search functionality),	ferings.	332
287	and a method to describe orchestrated execu-	<b>Mediation</b> In an orchestration, mediation	333
288	tion workflows through BPEL-WS.	can be done manually, yet facilitated by the	334
289	<b>Invocation</b> SOAP (Gudgin et al., 2007) is	inclusion of corresponding XML standards in	335
290	a protocol for remote procedure calls that	BPEL, e.g. functionality to transform vari-	336
291	are tunnelled through other protocols, most	ables (using XSLT transformations) and value	337
292	prominently HTTP POST requests, with the	selection (using XPath).	338
293	service endpoint as target URI of the request.	Around the core WS-* standards, a plethora	339
294	The content of such a request is an XML doc-	of other standards emerged, such as WS-	340
295	ument containing a SOAP envelope element,	Security, which introduced security features,	341
296	which is subdivided into header and body el-	WS-BPEL4People, which introduced human	342
297	ements, to describe the remote procedure call.	tasks into service compositions, and WS-	343
298	<b>Description</b> The WSDL (Web Services De-	Policy, which allows to describe requirements	344
299	scription Language) is an XML standard to	of a service regarding, e.g. security, quality of	345
300	describe service operations, including message	service, etc.	346
301	formats, typically in terms of XML schema for-	Overall, the Web services stack provided	347
302	mat descriptions, and bindings to a specific	very limited support for automation by agents,	348
303	protocol, essentially wrapping remote proce-	and was rather focused on enabling human	349
304	dures into XML. As such, the descriptions	developers to operate, compose, and invoke	350
305	were, despite human readable comments as	SOAP/XML-based services, and executing	351
306	part of the description, mostly focused on en-	pre-defined orchestrated executions.	352
307	listing API operations, and – borrowing from	<b>3.2 Semantic Web Services</b>	353
		<b>Technologies</b>	354
		In an attempt to overcome these deficiencies,	355
		Semantic Web Services attempted to enrich	356

357 Service descriptions, filling in the blanks, while  
358 still building on the Web Services ideas and  
359 base technologies.

360 **Description** Frameworks like WSMO (Ro-  
361 man et al., 2005), and OWL-S (Martin et al.,  
362 2005) are mainly semantic extensions of ser-  
363 vice descriptions, annotating overall capabili-  
364 ties, as well as inputs, outputs, pre-conditions,  
365 and effects (IOPEs), with declaratively de-  
366 scribed concepts (e.g using OWL or F-Logic),  
367 in order to facilitate ontology-mediated com-  
368 position by leveraging ontological reasoning to  
369 match (by subsumption) descriptions of re-  
370 quester goals to service capabilities. On top,  
371 the description can include, e.g. information  
372 about how a service makes use of other ser-  
373 vices in a choreography.

374 **Mediation** Especially WSMO put media-  
375 tion into the focus and defined mediation com-  
376 ponents, which are modular and reusable (OO,  
377 GG, WG, and WW mediators) to connect and  
378 reconcile differences between ontologies, goals,  
379 and Web services. In OWL-S, mediators are  
380 just yet another service (Paolucci et al., 2004).

381 Overall, Semantic Web Services remained  
382 largely an academic endeavour, partially be-  
383 cause service providers hardly supplied the  
384 required formal and unambiguous annota-  
385 tions: this shortage of real-world service de-  
386 scriptions kept repositories and engines—such  
387 as WSMO’s reference execution architecture  
388 (WSMX) (Haller et al., 2005)—at a mostly con-  
389 ceptual or prototypical level.

### 390 3.3 Summary

391 Both traditional Web Services and Semantic  
392 Web Services technologies share that they did  
393 not go significantly beyond the simple publish-  
394 find-invoke pattern. The addition of workflow-  
395 based orchestrations here merely added hard-  
396 wired processes, in the case of BPEL-WS,  
397 as well as prototypically leveraging declara-  
398 tive AI Planning techniques for composition,  
399 e.g. (Pistore et al., 2004).

400 Since then, XML-based messaging has been  
401 largely replaced by JSON as the standard  
402 data format, and SOAP has been replaced by  
403 RESTful interaction via HTTP, where discov-  
404 ery and composition is facilitated using hyper-  
405 media, instead of registries.

In summary, (Semantic) Web Services re-  
search was arguably focused on the automa-  
tion of discovery, mediation and composition  
from a requester-provider point of view.

## 410 4 Agentic AI

411 We argue that the recent advent of Agentic AI  
412 protocols (Anthropic’s Model Context Proto-  
413 col (MCP) and Google’s Agent2Agent Proto-  
414 col (A2A) as the main contenders) resembles  
415 many of the promises of (Semantic) Web Ser-  
416 vices. We thus outline that these “agentic”  
417 extensions of LLMs indeed solve main chal-  
418 lenges, but also show strong resemblance with  
419 (Semantic) Web Services.

### 420 Model Context Protocol

421 MCP enables LLMs to call and coordinate  
422 tools where an LLM (Host) acts as orchestra-  
423 tor, who knows MCP clients that connect to  
424 MCP servers accessible through JSON-RPC.

425 **Descriptions** MCP servers are described in  
426 natural language (typically explaining the ser-  
427 vice operation capability in human-readable  
428 terms), along with a template RPC call, il-  
429 lustrating the interface.

430 **Discovery** In MCP, discovery is facilitated  
431 typically by system prompt, where the descrip-  
432 tions are injected. At the same time, we note  
433 that dynamic discovery through a registry can  
434 be easily added to the MCP architecture. We  
435 illustrate this in Figure 4: to this end, we have  
436 implemented a prototypical architecture, i.e.,  
437 an MCP host using a “repository” MCP server  
438 that in turn provides access to a repository of  
439 other MCP servers dynamically. As a demon-  
440 stration use case, we provide a travel booking  
441 use case and resp. MCP servers through an  
442 anonymized GitHub repository<sup>2</sup>.

443 **Composition, Execution** When the LLM  
444 is triggered with a prompt, that can be un-  
445 derstood of a natural language described user  
446 *goal*. The LLM can choose to either re-  
447 spond directly or to instantiate one of its  
448 known MCP clients to solicit an answer for  
449 the prompt. When several MCP client calls  
450 are required, this is achieved through, typi-

<sup>2</sup><https://anonymous.4open.science/r/agentic-ai-implementations-DOC2/>




Figure 4: The Publish-Find-Invoke Pattern in MCP.

451 cally ReACT-style (Yao et al., 2023), repeated  
 452 decomposition of the goal.

453 **Mediation** The MCP host LLM may adapt  
 454 calls to the MCP servers through LLM-  
 455 mediated binding, leveraging the description  
 456 template, and background knowledge.

457 We thus argue that (i) MCP can be easily  
 458 extended to implement the above-mentioned  
 459 traditional WS publish-find-invoke pattern  
 460 and (ii) solves mediation/binding, by means  
 461 of relatively lightweight, natural language ca-  
 462 pability descriptions and the increasing ability  
 463 of LLMs to interpret those, and (iii) provides  
 464 dynamic orchestration through ReAct.

465 *Potential and Limitations.* Our prototypi-  
 466 cal tests in using MCP for a typical service  
 467 orchestration use case, conference travel book-  
 468 ing, confirms promising results, but also shows  
 469 that user control to when and which services  
 470 are called on behalf of a user goal is still lim-  
 471 ited, if we leave control to LLMs. Remarkably,  
 472 we can also show that dynamic enrichment of  
 473 services published on a dynamic repository is  
 474 entirely possible, but leaves questions such as  
 475 reliability of advertisements open: repsective  
 476 considerations from earlier initiative such as  
 477 UDDI, or “app stores” to ensure trustworthi-  
 478 ness of services on such repositories, should  
 479 probably be considered.

480 Also, the delegation of composi-  
 481 tion/mediation to an LLM may be considered  
 482 less reliable than the symbolic reasoning and  
 483 planning techniques envisioned in Semantic  
 484 Web Services. A final solution to this  
 485 challenge seems to call for a neurosymbolic  
 486 approach (Sheth et al., 2023; De Smet and  
 487 De Raedt, 2025).

488 We also note that so far, composition is, in  
 489 most documented MCP use cases, limited or-  
 490 chestration, where a single LLM takes the role  
 491 of an orchestrator of services that are known  
 492 at design time.

## 493 A2A

494 A2A by Google was presented as a comple-  
 495 mentary extension to MCP. A2A does not per-  
 496 se make the same strict client-server distinc-  
 497 tion as MCP, and also allows for HTTP-based  
 498 interaction. On top, A2A supports *discov-*  
 499 *ery* by so called AgentCards, which, as op-  
 500 posed to MCP’s simple text-based *descrip-*  
 501 *tions*, are more structured (e. g. making capa-  
 502 bilities explicit). Also, A2A uses unique IDs  
 503 for contexts and tasks, which can be shared  
 504 along with calls, and –similar to session IDs in  
 505 Web applications– enables (host) applications  
 506 to incorporate communication state informa-  
 507 tion into compositions, without revealing full  
 508 internal state details. Additionally, A2A sup-  
 509 ports some form of authentication and access  
 510 control. As such, A2A adds “convenience”  
 511 and conceptual promises on top of MCP that,  
 512 however, yet have to be proven in practical  
 513 use cases, for instance in terms of supporting  
 514 real choreographed, collaborative agent com-  
 515 positions.

516 However, it seems that –partially due to the  
 517 novelty of A2A– fully controllable mechanisms  
 518 or protocols to govern collaborative environ-  
 519 ments and adding safeguards for secure collab-  
 520 oration of agents, do not seem to be built-in  
 521 or existing as of yet.

## 522 5 Related Work

523 Together with the recent rise of Agentic AI,  
 524 papers about Agentic AI are written. Some  
 525 make economic arguments about the incen-  
 526 tives in Agentic AI and the Web (Yang et al.,  
 527 2025), others write about the notion of agency  
 528 in Agentic AI from a multi-agent systems per-  
 529 spective (Botti, 2025). Closest to our work,  
 530 (Petrova et al., 2025) see Agentic AI in the  
 531 history of FIPA agents, Semantic Web, and  
 532 multi-agent systems. Yet, for Semantic Web,  
 533 the paper stays on the vision level, and does  
 534 not have a conceptual Semantic Web Services

535  
536  
537  
538  
539  
540  
541  
542  
543  
544  
545  
546  
547  
548  
549  
550  
551  
552  
553  
554  
555  
556  
557  
558  
559  
560  
561  
562  
563  
564  
565  
566  
567  
568  
569  
570  
571  
572  
573  
574  
575  
576  
577  
578  
579  
580  
581  
582

treatment that we provide.

## 6 Outlook and Challenges

In summary, we believe one can learn important lessons from revisiting traditional Web Services terminology in the light of the current Agentic AI “hype”: while the main ideas and principles behind MCP or A2A based agents are not fundamentally new, LLMs and their reasoning capabilities have enabled automating mediation, binding and composition, in a manner and scale that has not been achievable by (Semantic) Web Services. However, some main challenges arise or remain:

- **agency:** i.e., what to delegate to whom? Sensitive aspects to achieve an agent’s goal may (not) be delegated to remote agents. In particular, fully-observable *functions* within the agent’s own capabilities *could* be executed locally, whereas service operations that require capabilities (i.e. access to the internal state) of another agent *require* remote operations.
- **collaboration:** each agent (and also service), may have own goal(s), i.e., the perspective of a single user agent defining the goal of a composed service execution is over-simplifying. In particular, *orchestration* is only one view of service composition, and which is coordinated by a single agent’s goal, where as complex, non-hierarchical interactions of multiple agents need *choreography* by different means of governance.
- **contracting:** non-functional aspects and “service” delivery (in the sense of a business value or product) is not fully observable and may need a-priori negotiations of contracts, as well as decentralized error-handling and dynamic contract renegotiation during execution. Especially, dynamic discovery and execution needs to cater for timeouts or temporal validity of service provision agreements, and may require guaranteed, standardized service-level agreements.
- **mediation:** in current agentic AI systems, mediation between the terminology of the different tools called seems to

be implicitly done by the LLM. The inner workings of this mediation, how it works and when it fails, next to the necessary negotiation in case something fails, deserves future study. On top, WSMO had identified other types of mediation, namely between goals and between the different communication paradigms and directions of services. Whether and how those more involved types of mediation can be addressed by LLMs is also an open question.

- **Web:** most digital information and functionality is currently provided via the Web, which is built as a resource-oriented and not a service-oriented architecture, and where hypermedia facilitates discovery and composition, all of which is yet to be leveraged by Agentic AI.

While we hardly see these challenges addressed in the current hype yet, they call for action and revisiting potentially ahead-of-their-time solution proposals from the research community.

## Limitations

- The lack of formal treatment in the Agentic AI proposed standards makes it hard to objectively tell what is the technical essence.
- Similarly, the Agentic AI proposed standards rarely state which architectural decisions have led to their development, which makes it hard to discuss the standards. This was also what motivated us to write the paper at hand, to investigate the intellectual roots of Agentic AI.

## References

Marco Aiello. 2018. *The Web Was Done by Amateurs - A Reflection on One of the Largest Collective Systems Ever Engineered*. Springer.

Gustavo Alonso, Fabio Casati, Harumi Kuno, and Vijay Machiraju. 2004. *Web Services: Concepts, Architectures and Applications*. Springer.

Alexandre Alves, Assaf Arkin, Sid Askary, Charlton Barreto, Ben Bloch, Francisco Curbera, Mark Ford, Yaron Goland, Alejandro Guizar, Neelakantan Kartha, Canyang Kevin

583  
584  
585  
586  
587  
588  
589  
590  
591  
592  
593  
594  
595  
596  
597  
598  
599  
600  
601  
602  
603  
604  
605  
606  
607  
608  
609  
610  
611  
612  
613  
614  
615  
616  
617  
618  
619  
620  
621  
622  
623  
624  
625  
626  
627  
628

629	Liu, Rania Khalaf, Dieter König, Mike Marin, Vinkesh Mehta, Satish Thatte, Danny van der Rijn, Prasad Yendluri, and Alex Yiu. 2007. Web services business process execution language version 2.0. Oasis standard. <a href="https://docs.oasis-open.org/wsbpel/2.0/OS/wsbpel-v2.0-OS.html">https://docs.oasis-open.org/wsbpel/2.0/OS/wsbpel-v2.0-OS.html</a> .	Jacek Kopecký, Tomas Vitvar, Carine Bournez, and Joel Farrell. 2007. <a href="#">SawSDL: Semantic annotations for wsdl and xml schema</a> . <i>IEEE Internet Computing</i> , 11(6):60–67.	683 684 685 686
636	Anthropic. 2025. <a href="#">Model context protocol specification, version 2025-06-18</a> . MCP Specification. Accessed: 2025-09-14.	David Martin, Massimo Paolucci, Sheila McIlraith, Mark Burstein, Drew McDermott, Deborah McGuinness, Bijan Parsia, Terry Payne, Marta Sabou, Monika Solanki, Naveen Srinivasan, and Katia Sycara. 2005. Bringing semantics to web services: The owl-s approach. In <i>Semantic Web Services and Web Process Composition</i> , pages 26–42, Berlin, Heidelberg. Springer	687 688 689 690 691 692 693 694 695
639	Tim Berners-Lee, James Hendler, and Ora Lassila. 2023. <i>The Semantic Web: A New Form of Web Content that is Meaningful to Computers will Unleash a Revolution of New Possibilities</i> , 1 edition, page 91–103. Association for Computing Machinery, New York, NY, USA.	Massimo Paolucci, Naveen Srinivasan, and Katia P. Sycara. 2004. <a href="#">Expressing WSMO mediators in OWL-S</a> . In <i>Proceedings of the Workshop on Semantic Web Services (SWS) at the 3rd International Semantic Web Conference (ISWC)</i> , volume 119 of <i>CEUR Workshop Proceedings</i> . CEUR-WS.org.	696 697 698 699 700 701 702
645	V. Botti. 2025. <a href="#">Agentic AI and multiagentic: Are we reinventing the wheel?</a> <i>CoRR</i> , abs/2506.01463.	Mike P. Papazoglou and Willem-Jan van den Heuvel. 2007. <a href="#">Service oriented architectures: approaches, technologies and research issues</a> . <i>VLDB J.</i> , 16(3):389–415.	703 704 705 706
648	Lennert De Smet and Luc De Raedt. 2025. <a href="#">Defining neurosymbolic ai</a> . <i>arXiv preprint arXiv:2507.11127</i> .	Charles J. Petrie. 2016. <i>Web Service Composition</i> . Springer.	707 708
651	Mark Endrei, Jenny Ang, Ali Arsanjani, Sook Chua, Philippe Comte, Pål Krogdahl, Min Luo, and Tony Newling. 2004. <i>Patterns: service-oriented architecture and web services</i> . Redbooks. IBM Corporation.	Tatiana Petrova, Boris Bliznioukov, Aleksandr Puzikov, and Radu State. 2025. <a href="#">From semantic web and MAS to agentic AI: A unified narrative of the web of agents</a> . <i>CoRR</i> , abs/2507.10644.	709 710 711 712
656	Dieter Fensel, Holger Lausen, Axel Polleres, Jos De Bruijn, Michael Stollberg, Dumitru Roman, and John Domingue. 2006. <i>Enabling semantic web services</i> . Springer Science & Business Media.	Marco Pistore, Fabio Barbon, Piergiorgio Bertoli, Dmitry Shaparau, and Paolo Traverso. 2004. Planning and monitoring web service composition. In <i>International Conference on Artificial Intelligence: Methodology, Systems, and Applications</i> , pages 106–115. Springer.	713 714 715 716 717 718
661	Google and A2A Project Contributors. 2025. <a href="#">Agent2agent (a2a) protocol specification</a> . <a href="https://a2a-protocol.org/v0.3.0/specification/">https://a2a-protocol.org/v0.3.0/specification/</a> . Version 0.3.0.	Chris Preist. 2004. A conceptual architecture for semantic web services. In <i>The Semantic Web – ISWC 2004</i> , pages 395–409, Berlin, Heidelberg. Springer Berlin Heidelberg.	719 720 721 722
665	Martin Gudgin, Marc Hadley, Noah Mendelsohn, Jean-Jacques Moreau, Henrik Frystyk Nielsen, Anish Karmarkar, and Yves Lafon. 2007. Soap version 1.2 part 1: Messaging framework (second edition). W3c recommendation. <a href="https://www.w3.org/TR/soap12/">https://www.w3.org/TR/soap12/</a> .	Dumitru Roman, Uwe Keller, Holger Lausen, Jos De Bruijn, Rubén Lara, Michael Stollberg, Axel Polleres, Cristina Feier, Christoph Bussler, and Dieter Fensel. 2005. Web service modeling ontology. <i>Applied ontology</i> , 1(1):77–106.	723 724 725 726 727
671	A. Haller, E. Cimpian, A. Mocan, E. Oren, and C. Bussler. 2005. <a href="#">Wsmx - a semantic service-oriented architecture</a> . In <i>IEEE International Conference on Web Services (ICWS'05)</i> , pages 321–328 vol.1.	Timo Schick, Jane Dwivedi-Yu, Roberto Dessí, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. Toolformer: language models can teach themselves to use tools. In <i>Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS '23</i> , Red Hook, NY, USA. Curran Associates Inc.	728 729 730 731 732 733 734 735 736
676	Timotheus Kampik, Adnane Mansour, Olivier Boissier, Sabrina Kirrane, Julian Padget, Terry R. Payne, Mumindar P. Singh, Valentina Tamma, and Antoine Zimmermann. 2022. <a href="#">Governance of autonomous agents on the web: Challenges and opportunities</a> . <i>ACM Trans. Internet Technol.</i> , 22(4).	Amit Sheth, Kaushik Roy, and Manas Gaur. 2023. <a href="#">Neurosymbolic artificial intelligence (why, what,</a>	737 738

739 and how). *IEEE Intelligent Systems*, 38(3):56–  
740 62.

741 Ashish Vaswani, Noam Shazeer, Niki Parmar,  
742 Jakob Uszkoreit, Llion Jones, Aidan N. Gomez,  
743 Łukasz Kaiser, and Illia Polosukhin. 2017. At-  
744 tention is all you need. In *Proceedings of the 31st*  
745 *International Conference on Neural Information*  
746 *Processing Systems*, NIPS’17, page 6000–6010,  
747 Red Hook, NY, USA. Curran Associates Inc.

748 Yingxuan Yang, Mulei Ma, Yuxuan Huang, Hua-  
749 can Chai, Chenyu Gong, Haoran Geng, Yuan-  
750 jian Zhou, Ying Wen, Meng Fang, Muhao Chen,  
751 Shangding Gu, Ming Jin, Costas J. Spanos,  
752 Yang Yang, Pieter Abbeel, Dawn Song, Weinan  
753 Zhang, and Jun Wang. 2025. [Agentic web:](#)  
754 [Weaving the next web with AI agents.](#) *CoRR*,  
755 abs/2507.21206.

756 Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak  
757 Shafran, Karthik Narasimhan, and Yuan Cao.  
758 2023. React: Synergizing reasoning and acting  
759 in language models. In *International Conference*  
760 *on Learning Representations (ICLR)*.