Journal of Ethics and Legal Technologies – Volume 8(2), November 2026

AI agency language without moral personhood: a two-register framework for responsibility attribution

Indrajith P. Karunanayaka

Independent Researcher, Gampaha, Sri Lanka

Abstract: Artificial intelligence systems are described as deciding, rejecting, recommending, apologising, and accepting responsibility, and these descriptions are written into documentation, interfaces, organisational records, policy texts, and reporting. This article treats agency terminology as a problem of description and develops a framework that separates functional description from morally weighty attribution. Because such terminology moves across technical, institutional, legal, and public contexts, unmarked shifts in meaning can affect where responsibility is located. Using comparative conceptual reconstruction together with a repeatable analytic procedure, the article reconstructs the distinction in Theravāda Abhidhamma between *sammuti*, or conventional designation, and *paramattha*, or ultimate moral-analytic description. It then applies the framework to conceptual test cases involving conversational systems, automated institutional decisions, and AI-mediated information services. The analysis distinguishes an operational-conventional use, a social-legal-conventional use, and an ultimate moral-analytic register. Unmarked shifts among them turn useful shorthands into moral attributions and thereby conceal the designers, organisations, and regulators responsible for system behaviour. Rather than eliminating agency language, it is therefore important to mark which use is in play. Recent work has recognised imprecise AI terminology as a problem in its own right, but has not explained how agency terms shift between descriptive registers. The article explains one important form of imprecision, provides a procedure for identifying it, and offers a protocol that maps each register to a responsibility structure so that documentation, interface texts, and governance rules can describe system behaviour without concealing who is accountable.

Keywords: Artificial intelligence, Responsibility attribution, Moral personhood, Theravāda Abhidhamma

1. Introduction

In product documentation, news reporting, policy drafting, and ordinary conversation, artificial intelligence (AI) systems are commonly described in the language of agency. A chatbot "decided" not to answer a sensitive question. An automated hiring tool "rejected" a candidate. A clinical decision-support system "recommended" a treatment. An autonomous vehicle "chose" to swerve. A model "apologised" for an oversight or "accepted responsibility" for a fault. Industry practice follows the same pattern. Families of systems built around large language models that plan on behalf of human principals, call tools, browse websites, and execute multi-step tasks are now named by the term "agentic AI" (Raptis et al., 2025).

This terminology is useful. It allows engineers to describe complex behaviour briefly, lets users interact with interfaces through familiar terms, helps reporters compress a story, and enables regulators to name an object of oversight. The difficulty is that these same words determine who is blamed. The way a failure is described shapes who is judged in connection with it, whom users trust, and where institutions and courts direct attention when asking who is responsible. When an operational shorthand becomes a claim that the system intended, chose, judged, or is responsible for something, the language stops describing behaviour and begins assigning moral responsibility. This is not only a social question about AI. It is also a question for information ethics and for the study of how systems are described and recorded, as well as for moral philosophy concerning agency and responsibility. The words used to describe a system distribute moral and legal accountability among the people and organisations that build, deploy, and control it.

Many of these descriptions are also written down. A model card records what a system does under stated conditions. An interface states what the system tells a user. A rejection letter sets out what was decided about an applicant. A regulatory filing captures the state of a system before an incident, and a news report recounts what an organisation said afterward. Such records do more than record behaviour. Depending on their wording, they establish the conditions under which a later reader understands what acted, what was intended, and who should be held responsible for the outcome. These records also outlast the events they describe. An auditor, a journalist, a court, or a later user will consult the document rather than witness the moment it describes.

Agency language is too useful to abandon and too consequential to leave undisciplined. Rejecting all agency talk would deprive ethics, law, design, and ordinary speech of the functional terminology needed to describe what

these systems do. Accepting it without discipline leads to four forms of slippage. Ontological slippage treats the phrase "AI decided" as evidence about what it is like to be a system, namely that something resembling consciousness or sentient volition exists. Moral slippage treats the system as a bearer of praise or blame for its outputs. Responsibility slippage lets design, deployment, and control decisions made by humans disappear behind the system. Regulatory slippage carries this ambiguity into policy and law, producing overreach when systems are treated as quasi-persons and underreach when accountability is weakly located.

Each of these problems has already been examined. Work on anthropomorphism explains why people readily attribute mental states to artificial systems (Coeckelbergh, 2010; Damiano and Dumouchel, 2018). The responsibility-gap literature beginning with Matthias (2004) and Sparrow (2007) develops a typology of culpability and accountability for autonomous machines (Santoni de Sio and Mecacci, 2021), although others argue that the gap has been overstated (Tigard, 2021). Analyses of role-responsibility and levels of abstraction show that thin, undemanding forms of agency or responsibility can be attributed to systems without treating them as moral persons (Floridi and Sanders, 2004; Vacek, 2026). Relational accounts ground moral standing in interaction rather than internal properties (Coeckelbergh, 2014; Gunkel, 2018). Critiques of legal personhood show how it can distance accountability from human and institutional actors (Bryson et al., 2017). Information science has begun to treat AI terminology as an object of study in its own right, arguing that it is inconsistent and often misleading (Kozak, 2026). None of this work yet supplies a register-based analysis of the language. Such an analysis would specify when agency talk can be accepted as shorthand, when it becomes a category mistake, and what responsibility structure each use implies.

This article develops such an analysis. It reconstructs the distinction in Theravāda Abhidhamma, a tradition of analytic moral psychology, between *sammuti*, or conventional designation, and *paramattha*, or ultimate moral-analytic description. It then uses that distinction to build a Two-Register AI Agency Language Framework. The framework legitimates operational and social-legal agency talk within their registers while blocking unmarked drift toward claims that current systems possess sentient volition, moral personhood, or responsibility-bearing subjectivity. The reconstruction uses the relevant tradition as an analytic resource for analysing how artificial systems are described across technical, institutional, and governance contexts. The argument is restricted to current and near-future AI systems.

The article first sets out its methodology and contribution. It then examines how technical, interface, and commercial uses of agency language

merge, how that merging affects responsibility attribution, and where existing work in AI ethics and information science leaves a gap. The following sections reconstruct the conventional and ultimate moral-analytic registers, develop the framework, and apply it to three focal cases and a short language protocol, after which objections are considered and implications for documentation, governance, and accountability are drawn.

2. Methodology and Contribution

This article is a conceptual study. Its methodology is based on conceptual analysis within information ethics and applied moral philosophy, and its object of study is the way AI systems are described in documentation, interfaces, institutional records, policy texts, and reporting. It does not report empirical findings, and its examples are illustrative conceptual test cases rather than data. The methodology has two parts. The first is comparative conceptual reconstruction, in which a distinction drawn from one tradition of moral psychology is restated in analytic terms and examined for its usefulness in a new domain. The second is a repeatable analytic procedure for investigating a given use of agency language.

The procedure consists of five steps. First, identify an agency phrase that occurs in use, such as "AI decided," "the system rejected the candidate," or "I take responsibility." Second, determine which register the phrase occupies, whether it describes the operational behaviour of an artefact, names an object of governance and law, or makes a claim about consciousness, volition, and moral standing. Third, examine whether the wording encourages a shift from functional description to the attribution of moral personhood or responsibility-bearing subjectivity, and whether that transition is marked or passes unnoticed. Fourth, when such drift occurs in a context with moral or legal stakes, redescribe the situation in language that preserves accountability. Fifth, locate responsibility in the relevant human and institutional actors. The procedure is designed for written or published descriptions where the stakes of attribution are highest. Ethicists, designers, technical writers, regulators, and reporters can all apply it in the same way, and the test cases below show how it operates.

The conceptual resource for this procedure is the distinction in Theravāda Abhidhamma between conventional and ultimate description. Abhidhamma is used here for three analytic reasons. It supplies a fine-grained distinction between conventional designation and ultimate moral-analytic description, and it draws that distinction at the level of how a situation is described rather than at the level of metaphysics or relations. It also provides a clear criterion for morally weighty agency, namely the presence

of conscious mental events with the appropriate volitional structure. This gives the framework a clear standard for what must exist in a system before moral attributions are warranted. In addition, it fits the agency-language dilemma closely because the tradition is built around avoiding the same kind of category mistake that the dilemma involves. This reconstruction is selective, and it uses only those elements of Abhidhamma needed to distinguish conventional description from morally weighty agency.

The contribution should be distinguished from neighbouring claims. The article does not argue merely that AI lacks consciousness, that anthropomorphism is dangerous, or that responsibility gaps exist, nor does it deny any of these. Its contribution is a register-based discipline for describing AI systems and attributing responsibility wherever system behaviour is documented, presented to users, or entered into an institutional record. The discipline legitimates operational agency talk within its proper register, prevents drift toward moral personhood, sentient volition, or responsibility-bearing subjectivity, and maps each register to a specific responsibility structure. Previous work identifies the risks of agency language or provides positive accounts at the level of moral or legal status. This framework intervenes at the level of description and ties that intervention directly to the location of accountability.

3. Agency Language and Responsibility Attribution

3.1 How the Registers Merge

Three factors have contributed to the widespread use of agency language in AI discourse. The first is the technical history of the term agent. In standard textbook use, an agent is anything that perceives an environment through sensors and acts on it through actuators (Russell and Norvig, 2010, 34). This definition is functional. It says nothing about consciousness or moral status, and the same holds for the term agent in multi-agent systems and reinforcement learning. The second factor is product and interface design. From early conversational systems to current language model assistants, designers have used first-person address, personal names, voices, and conversational style to make systems easier to use. Giving a system a personal name, as with widely used voice assistants, encourages users to perceive it as an entity that listens, acts, and has a mind of its own (Lind, 2025). The third factor is commercial. With the rise of systems that combine planning, tool use, and memory, the phrase agentic AI has moved out of the

research vocabulary and into the marketing of products described as agents that handle workflows, carry out transactions, and represent users.

These three uses reinforce one another. The technical meaning is functional, the interface meaning is performative, and the commercial meaning relies on users' tendency to read mental states into human-like outputs. Once they operate together, a sentence such as "AI decided" carries functional, interface, and quasi-moral meanings at once, and readers are rarely equipped to distinguish them. Critical work on corporate communication calls the strategic or careless use of such human-like presentations human-washing and locates it within a broader pattern of business framing that misleads stakeholders about what a system can actually do (Scorici et al., 2022). Corpus studies of how voice assistants and chatbots are described in the media show that the persistent attribution of agency, animacy, and intention to artificial systems creates a stable cultural picture of them as human-like actors (Lind, 2025). What appears to be neutral word choice is, in these settings, a way of assigning agency, and it changes what users expect, whom they trust, and whom they blame.

3.2 Effects on Responsibility Attribution

The most important of these effects concerns the attribution of responsibility. When an automated system fails, the language of distributed agency can create what Elish (2019) calls a moral crumple zone, a situation in which the human nearest the interface absorbs the moral and legal costs of a failure that person did not control. This pattern follows the wording. If a description states that an agent acted and the institutional and design context remains invisible, the visible human becomes that agent. Incident reports and policy frameworks then repeat this misattribution. The responsibility-gap literature describes a complementary pattern. In the account offered by Matthias (2004), learning systems generate consequences no individual can fully predict or control, opening a gap in moral attribution. That argument is extended to lethal autonomous weapons by Sparrow (2007). Several kinds of gaps are distinguished by Santoni de Sio and Mecacci (2021), while according to Tigard (2021), no genuine techno-responsibility gap exists and the resources of moral philosophy already permit responsibility to be traced across complex systems. More recent work pursues the same aim from the constructive side, specifying conditions for assigning responsibility even where an outcome appears uncontrolled, such as to those who voluntarily initiated an automated process and expected benefits from it (Gong, 2026). That dispute need not be resolved here. What matters is how agency language interacts with whichever conceptual structure is used for assigning

responsibility. If the language is left undisciplined, even a strong structure becomes harder to apply because the words themselves have already moved the apparent locus of agency away from the people and institutions involved.

The phrase "AI decided" is not always wrong. In an engineering meeting it may be the most efficient way of stating that a model produced a particular output under specified conditions. In a courtroom, a regulatory filing, a news report, or a corporate apology, the same sentence carries very different stakes, and without a disciplined framework it slips between these settings without any signal that something has changed. Written descriptions make this situation harder to recognise because a sentence drafted for one audience is often reused by another. Before any problem of responsibility attribution arises, therefore, the problem begins with wording.

3.3 Existing Responses and the Remaining Gap

Several bodies of literature respond to parts of this problem. Work on anthropomorphism explains why humans attribute mental states to artificial systems and why such attributions can mislead, whether they are treated as a constitutive feature of social robotics or as a source of distorted expectations about a system's capabilities (Coeckelbergh, 2010; Damiano and Dumouchel, 2018; Sharkey and Sharkey, 2012). This work is largely diagnostic and does not say when agency language is acceptable. A second stream offers positive but status-level accounts. According to Floridi and Sanders (2004), at an appropriate level of abstraction a system can be considered a moral agent in a thin sense even without consciousness, while Johnson (2006) warns against slipping from such functional descriptions to moral or ontological claims. Through role responsibility, Vacek (2026) attaches a philosophically undemanding form of responsibility to AI systems without treating them as moral persons. Relational accounts ground moral standing in interaction rather than internal properties (Coeckelbergh, 2014; Gunkel, 2018). Critiques of legal personhood warn that granting it to synthetic entities is morally unnecessary and can shield producers and deployers from accountability (Avila Negri, 2021; Bryson et al., 2017; Chopra and White, 2011). Intercultural work argues that AI ethics should draw on non-Western resources as well (Hongladarom, 2020; Hongladarom and Bandasak, 2023; Lin, 2023). More directly, Soulier (2026) frames machine agency as a problem of conceptual extension. The relevant question is not whether a system fits a definition of agency. It concerns what normative function the application of the concept of agency to machines performs. This article shares that normative orientation and its concern with responsibility but locates

the problem differently. The question here is not whether agency should be extended to machines, but how a single agency term moves among operational, social-legal, and moral-analytic uses and how that movement shapes responsibility attribution.

A further strand of relevant work examines AI terminology and informational objects from within information science. Kozak (2026) argues that AI terminology is inconsistent and often misleading, illustrating the point through the terms AI peer review and AI ethics and calling on the research community to reconcile its vocabulary. Hjørland (2026) raises a connected issue about a long-standing misnomer in the field, arguing that treating document retrieval as information retrieval has encouraged unrealistic goals and that generative systems have intensified rather than resolved this confusion. In parallel with this work on terms, information science has begun to ask what AI outputs are as informational objects. Donner (2025) examines whether large language model outputs can be considered documents and introduces the concept of artificially blended testimony to place them within documentation activity. Prompting is treated by Thellefsen et al. (2026) as a communicative and epistemic act rather than a technical input. According to Senn (2026), mediating systems, including AI systems, provide the conditions under which a reader can know of that act without participating in it. Taken together, these works show that AI terminology and AI-generated informational objects have become explicit objects of conceptual analysis. None of them yet explain the relevant mechanism. Showing that a term is ambiguous is not yet an analysis of how a single term changes meaning as it moves from engineering description to institutional record and then to moral claim.

The risks of agency language are widely identified in the ethics literature, but a disciplined positive account of its use is not offered. Positive accounts, including role responsibility, levels of abstraction, and relational standing, are located at the level of moral or legal status rather than at the level of language. Across these literatures, AI terminology remains unsettled, but existing work does not explain how a single agency term changes between operational, social-legal, and moral-analytic uses. The remaining gap is a register-based framework that legitimates operational agency talk within its proper register, prevents drift toward moral-ontological reification, and maps each register to a specific responsibility structure. Because agency descriptions are written into documentation, interfaces, institutional records, governance texts, and public communication, such a framework can be applied at the site where those descriptions are produced, as indicated in Section 8. The next two sections build it, beginning with the conceptual distinction on which the framework rests.

4. Conventional and Ultimate Moral-Analytic Registers

4.1 Abhidhamma as Analytic Moral Psychology

Theravāda Abhidhamma is an extensive tradition of analytic moral psychology set out in canonical texts as well as in commentarial and handbook literature (Bodhi, 2007; Karunadasa, 2016). Although it developed within Theravāda Buddhism, it also functions as a technical analysis of mental events, mental factors, and morally significant action. This article uses it as a conceptual and moral-psychological resource. The borrowing is selective and limited to terms directly relevant to agency language. The resulting framework should be judged on an analytic basis alongside contemporary phenomenological, relational, and levels-of-abstraction approaches.

4.2 *Sammuti* and *Paramattha*

The central distinction lies between *sammuti*, or conventional designation, also called ordinary-language truth, and *paramattha*, or ultimate moral-analytic description. As Karunadasa (2016, 68-70) explains, these two terms do not name two grades of truth, one lower and one higher. They name two ways of presenting what truth is, both valid in their proper register, and conflating them generates category mistakes. In ordinary usage, descriptions of persons, possessions, and social roles operate in the conventional register. Saying that a named person is walking is true at the level of conventional designation, even though according to Abhidhamma analysis that name is a conceptual designation rather than an enduring substance and is ultimately described as a stream of arising and ceasing mental and physical processes. The tradition does not treat conventional description as a falsehood to be discarded. It regards that description as a workable description appropriate to one register, and it treats ultimate analysis as a more fine-grained description of the same situation.

This analytic move resembles several familiar Western approaches. Wittgenstein's concept of language games shows how the same words perform different roles in different settings (Wittgenstein, 2009, 13e, 15e). Searle's analysis of institutional facts indicates how social and legal practices can be maintained without further claims about underlying physical processes (Searle, 1995). Comparable lines are drawn by the distinctions between manifest and scientific images (Sellars, 1962) and between personal and sub-personal levels (Dennett, 2010). The Abhidhamma distinction is not unique to any of these, but the resemblance is close enough that non-Buddhist readers can approach it as a relevant analytic strategy.

Because the dilemma described above takes the same form, this structure is what makes the distinction useful for AI agency language. The phrase "AI decided" can be a useful conventional, operational description. When it is read as an ultimate moral or ontological claim, it becomes a category mistake. The distinction permits both points to be stated at once. Conventional agency language is acceptable within its register, and the same words become a category mistake when they are carried without marking into the register of moral and ontological analysis. Understood in this way, the distinction functions as a theory of descriptive registers. A single situation can be described in more than one correct way. Each description answers to its own criteria, and difficulty begins when a description made for one purpose is read as an answer to a different question. This capacity for register-relative description makes the distinction useful for studying how systems are described and how those descriptions are recorded and read.

4.3 A Criterion for Morally Weighty Agency

Three further terms make this application precise. *Citta* is consciousness, understood as a conscious mental event or a moment of awareness. *Cetasika* refers to the mental factors that arise together with *citta* and shape its character. One of these factors is *cetanā*, usually translated as volition or intention, which in this analysis coordinates the other factors with respect to an object and is regarded as the factor from which morally significant action arises (Bodhi, 2007, 26-29, 80; Nyanaponika Thera, 1976, 41-43, 57, 138). These terms fix what counts as morally weighty agency within the tradition. Morally significant action requires a stream of conscious mental events with the appropriate volitional structure. It is not enough for a system to produce outputs that resemble decisions, because the right kind of coordinated volitional event with the relevant factors must also occur within a stream of awareness.

This criterion is used here as a comparison rather than as a doctrine to be imported into AI ethics. Under it, attributing *citta*, *cetasika*, or *cetanā* in the relevant sense to current and near-future AI systems is not justified. Functional goal-directedness and sentient volition are different, and operational agency language can describe the former without committing to the latter. The distinction is useful for separating accountability regardless of whether one holds a Cartesian, physicalist, functionalist, or Abhidhammic account of mind, so nothing in the argument depends on the Abhidhamma framework.

One clarification is needed here. Abhidhamma also rejects an enduring self in humans and decomposes the person at the ultimate register into

streams of mental and physical processes (Gethin, 1998, 208-211). If humans are also no-self, why distinguish them from AI systems? The answer is that humans instantiate streams of conscious mental events with volition, which the no-self analysis does not deny. What it denies is that those streams together constitute a substantial enduring self. The relevant difference between humans and current systems is not the possession of a self, since the tradition attributes a self to neither, but the presence or absence of the right kind of conscious mental processes. That analysis also brings interdependence to the fore, and it suggests that responsibility should be distributed across the humans, datasets, and institutions involved rather than settling by linguistic default on the most visible artefact.

5. Two-Register AI Agency Language Framework

5.1 Two Registers, Three Functional Uses

The framework distinguishes two main analytic registers, the conventional and the ultimate moral-analytic. Within the conventional register it further distinguishes two uses, the operational and the social-legal. The framework's name reflects this primary division. In the sense most important for moral attribution, there are two registers, and the conventional register covers two distinct uses. Speakers may move among these uses, but they must signal which use a claim belongs to and avoid carrying a claim across them without warrant.

In operational-conventional use, claims describe AI systems as artefacts that take inputs, maintain internal state, and produce outputs that perform specific tasks. Even verbs such as predicts, selects, executes, navigates, negotiates, recommends, and decides can serve as functional shorthands here, and claims are assessed by functional and performance criteria rather than by attributing consciousness or moral status.

In social-legal-conventional use, claims describe systems as named objects of governance, regulation, contract, and public discourse. They permit sentences stating that an automated hiring system is subject to audit, that an autonomous vehicle is registered for trial use in a jurisdiction, or that a clinical system requires informed consent from the patient. Such claims do not commit the speaker to any view about the system's inner life. They are conventions established by institutions and law.

In the ultimate moral-analytic register, claims concern whether a system instantiates conscious mental events with the appropriate volitional structure, whether it bears moral responsibility, and whether it has the standing

of a moral person. Under the criteria given above, current and near-future systems do not warrant attributions of sentient intention, moral responsibility, or moral personhood within this register.

These are registers of description rather than document types. A model card typically operates in the operational use, a compliance policy or regulatory filing in the social-legal use, and a sentence reporting that a system intended an outcome crosses into the ultimate register wherever it appears. A single document can move among uses within one paragraph, and that is one reason the movement is easy to miss.

The framework does not prohibit agency language. Operational and social-legal uses of terms such as agent, decides, and similar expressions are legitimate within their registers. What it blocks is the unmarked movement of such claims into the ultimate moral-analytic register. Three movements in particular are blocked. The first treats "AI decided" as evidence of sentient volition. The second treats "AI is responsible" in social-legal use as evidence of moral responsibility in the full sense. The third treats "AI is an agent" in any conventional use as a licence for the conclusion that the system is a moral person. Each is a category mistake of the kind identified by the *sammuti-paramattha* distinction. The framework also permits a positive step. A speaker can hold at once that operational agency language is appropriate for describing what a system does and that, at the ultimate register, the system is not a moral person or a bearer of responsibility. These two claims occupy different registers and do not conflict.

5.2 Functional Intention and Sentient Volition

Current agentic systems plan, set sub-goals, and pursue them through tool calls, and this raises an obvious objection. Why deny intention to them? The answer distinguishes two meanings. Functional intention is system-level goal-directedness, described by the system's architecture and behaviour, and it belongs to operational-conventional use. Sentient volition is something different, in the sense of *cetanā* as a coordinating factor within a stream of conscious mental events, and it is what the ultimate register requires. Accepting functional intention as a useful operational concept does not commit one to attributing sentient volition, and denying sentient volition does not require denying that a system displays goal-directed behaviour. The distinction only separates the two meanings, and that separation is what the agency-language dilemma requires.

5.3 Register Tests for Unmarked Drift

A small set of diagnostic questions, concerned with a sentence's syntax and pragmatics rather than with mechanical rules, helps a speaker locate a use of agency language and identify drift between registers. Does the construction confuse operational output with ontological status, so that the same verb describes a model selection and a moral commitment without qualification? Even when only the operational meaning is intended, do verbs such as decides, intends, and apologises carry moral connotations that invite an inference to sentient volition where only functional selection is justified? Does the sentence place the system in the agent position, remove designers, vendors, deployers, or operators from view, and relocate responsibility away from the actors who shaped the outcome? Where does the sentence occur, in a legal or policy text, in a technical report on which others rely, or in a public-facing statement? If the setting requires a specific subject of responsibility, an unmarked operational use will introduce a category mistake into the record. Would a non-expert reader in the relevant audience understand the phrase as attributing consciousness, intention, remorse, blame, or responsibility to the system? And is there a functional alternative that preserves accountability and should be preferred in high-stakes settings? These tests apply across categories of systems, but the drift is more acute for chatbots with fluent first-person language that invites such drift more readily and for systems that operate in the physical world.

6. Conceptual Test Cases and a Language Protocol

Three focal cases are examined in depth, each selected because it reveals a different form of slippage and a different responsibility structure. The first concerns what a conversational system says about itself. The second concerns what an organisation says about an automated decision. The third concerns what a supplier says about an information service. In each case the same wording is read across registers, the moral cost of unmarked drift is identified, the actors concealed by the wording are named, and a register-disciplined redescription is proposed.

6.1 Chatbot Apologies and First-Person Responsibility

A user encounters a sentence such as “Sorry, I cannot help with that request,” or more strikingly, “I should not have said that earlier, I apologise for it,” or “I take responsibility for the error.” In operational-conventional use the system has produced a string conditioned by its training and safety policies. In social-legal-conventional use the deploying organisation has chosen to present such strings to users. In the ultimate moral-analytic register the

system is not the kind of thing that can feel remorse or bear responsibility. The moral cost here comes from the first person. Because the system's outputs are not attempts to express something it believes or grasps (Bender et al., 2021), fluent first-person speech invites the user to impute to the system a communicative intent and an inner moral life that it does not have. An apology or an acceptance of responsibility then positions the system as a moral interlocutor, and the user looks to the system for accountability instead of looking to the organisation that deployed it. The concealed actors here are the developers and the deploying organisation, who remain responsible for the content, the quality, and any legal liability. The proposed redescription frames the incident as system behaviour rather than personal remorse, for example by reporting that the response was wrong and has been logged for review, and by stating the organisation's responsibility directly. Once the wording is corrected, the responsibility structure becomes visible because the organisation, not the artefact, is responsible for the error.

6.2 Automated Hiring and Decision Support

An automated screening tool rejects a candidate, who receives a message written in the language of decision and judgment. In operational-conventional use, a model has assigned a value below a threshold. In social-legal-conventional use, the hiring organisation is responsible for using the tool, including compliance with anti-discrimination law. In the ultimate moral-analytic register, the system does not reject the candidate in any morally weighty sense, because the system performs no judgment of the kind that moral appraisal presupposes. The moral cost here is that the language of judgment makes an institutional decision appear as a verdict delivered by the system itself. This obscures the organisation's accountability and leaves the candidate uncertain about what recourse is available. The concealed actor is the hiring organisation.

The redescription states that the application was screened out by an automated system used by the organisation under stated conditions, with specific review options. This aligns the wording with the responsibility structure, in which the organisation remains responsible for the lawfulness and fairness of the tool it chose to deploy. The same analysis extends to clinical decision support. A system that outputs a ranked list under specified inputs should not be described in terms that lead patients or junior clinicians to treat its suggestion as morally authoritative, because clinicians and institutions remain responsible for the decision and for informed consent.

6.3 AI-Mediated Information Services

A user asks a question and receives a synthesised answer instead of a list of sources. The answer itself, or the text around it, describes the system as something that knows and judges. It "knows" current guidelines, "discovered" that an earlier figure was wrong, "decided" that one source is more reliable than another, and "understands" the question. Descriptions of this kind now appear in product copy, in help documentation, in reference and discovery services that place a generative layer over a catalogue or database, and in reporting about what such systems can do.

In operational-conventional use, the system has retrieved, ranked, weighted, and generated text under specified inputs and model conditions. In social-legal-conventional use, a supplier or institution has selected the system, configured it, and presented its output as a service, and that actor is responsible for the reliability of what is offered and for the conditions under which it is offered. In the ultimate moral-analytic register, verbs such as knows, understands, and judges are not justified, because they report the presence of an epistemic subject and the system is not such a thing. Senn (2026) raises a related point from within information science, arguing that mediating systems, including AI systems, provide the conditions under which a reader can know of that act without participating in it.

The moral cost here differs from the cost in the previous two cases. What goes wrong first is a knowledge claim rather than a responsibility claim. Saying that a system found a source reliable presents a ranking step as an act of assessment. Saying that a system knows something presents generated text as testimony from a checked source. This invites the reader to grant the system the kind of trust historically given to a cognitive authority, in Wilson's sense of a source whose word carries weight (Wilson, 1983, 15). Hjørland (2026) supplies a structural reason for concern, since generative systems cannot reliably connect their answers to the sources on which they are based. Yet wording of this kind can imply that such a connection has been built. When the description hides that gap, the user cannot perform source criticism because the sources to be checked are not shown.

The concealed actors here are the supplier who built and tuned the system, the organisation that chose to place it before users, and the sources that provide the material on which the output depends. The proposed redescription keeps both the operation and the institution in view. It states that the system generated a summary from sources retrieved under stated conditions, points to those sources, and names the organisation responsible for the reliability of the service. The responsibility structure that then becomes visible runs across suppliers, deploying institutions, and underlying sources, each of which holds a specific place in how the answer was produced.

6.4 Systems Operating in the Physical World

The three focal cases above all concern settings in which the output is information. The framework is not limited to those. After a collision, a press release may state that "the vehicle decided to swerve," and the same sentence may appear in early news reports and regulatory filings. In operational-conventional use, a perception and control stack has produced an output under specific sensor and software conditions. In social-legal-conventional use, the manufacturer, the software supplier, the fleet operator, and the driver each hold a specific legal position. In the ultimate moral-analytic register, the vehicle is not a moral agent. The moral cost here is the cost documented by Elish (2019). Once the vocabulary allows the system's selection to be treated as a morally salient act, the available human involved in the process becomes a moral crumple zone, and the responsibility-gap literature shows how responsibility attribution then fails to reach anyone else (Matthias, 2004; Sparrow, 2007). Although physical control systems are structurally different from language models, public and regulatory discourse describes both through the same agency vocabulary, and the same redescription applies. Stating that the vehicle's control system produced the manoeuvre under the perceived conditions keeps the manufacturer, the software supplier, and the fleet operator in view, and it distributes responsibility among them rather than concentrating it on the vehicle or, by default, on the nearest human.

6.5 A Language Protocol

These cases share a structure that can be condensed into a short protocol. Table 1 collects a range of agency phrases, including extensions to enterprise agents, deceptive outputs, and claims about knowledge and understanding that are not examined in depth above. For each phrase it names the register the phrase confuses, explains the slippage, proposes a register-disciplined alternative, and states the responsibility implication. The protocol is a starting point rather than a finished tool, and specific sectors will require finer distinctions. Its structure is general enough to be applied directly in documentation, interface texts, and organisational communication.

The protocol does not require the elimination of conversational design. Friendly first-person address and a warm welcome often serve usability, accessibility, and trust, and they are legitimate in operational-conventional use. The question is not whether an interface may say "I can help" but whether, in contexts where social and legal stakes activate the ultimate register, it says or implies "I intended," "I apologise in a morally responsible way," or "I take responsibility."

7. Objections and Limitations

Several objections have been addressed along the way, and the strongest are considered here.

The framework underestimates what current systems do. Recent systems plan, call tools, manage memory, and act on environments, and Floridi and Sanders (2004) argue that at an appropriate level of abstraction they can be moral agents in a thin sense. The framework does not reject this. In line with Johnson's (2006) warning against drifting from functional descriptions to ontological claims, the approach can be read as a way of bringing the levels-of-abstraction analysis into the wording itself. It marks the level at which a thin moral-agent description is justified and does not carry such descriptions into the ultimate register without further argument.

Legal personhood is itself a convention, so why not extend it to AI systems? The framework does not block all uses within the social-legal register. It blocks slipping from legal convention to a claim about moral personhood at the ultimate register. As Bryson et al. (2017) argue, granting legal personhood to synthetic entities is morally unnecessary and can shield human and institutional actors from accountability. Making the register shift explicit strengthens that argument. The same applies to enterprise agents entering contracts, where legal capacity and liability depend on the contracting organisation rather than on the system.

Engineers should call these systems agents. The framework accepts that use in operational-conventional use. The only requirement is that engineering terms should not be carried into ethics, law, or public communication, where the stakes of attribution differ, without further argument.

This is merely a matter of terminology. The objection is partially correct, because nothing here depends on a hidden fact about machines that could be revealed by better wording. What it misses is that a description is not merely a record. It is the form in which what a system did becomes available to a reader who was not there, and it is also the form in which an organisation's interpretation of an event enters a record on which others act. The same phrase written in a technical document, a message to a customer, and a regulatory filing makes different actors visible in each and allows different questions to be asked. Terminology is what makes responsibility visible or keeps it obscure. That is why terminological confusion in AI is now treated as a problem for the research community rather than a matter of taste (Hjørland, 2026; Kozak, 2026).

Relational accounts already address the problem. These approaches hold that moral standing emerges through interaction (Coeckelbergh, 2014; Gunkel, 2018), and the framework is complementary to them. This literature ex-

plains whether and how artificial systems can have moral standing through interaction, while the framework asks how agency language should be used when the responsibility structure of current practice is considered. A relational claim can be stated in social-legal use without committing to ultimate moral-analytic attributions. Register discipline in turn helps prevent users' felt closeness to a system from being read as a claim about its legal or moral personhood.

Another objection concerns cultural specificity. If Abhidhamma belongs to a specific Buddhist tradition, why should it have authority in a debate outside that tradition? The framework does not depend on it as the only path to that insight. Related lines are drawn by the distinctions between manifest and scientific images (Sellars, 1962), brute and institutional facts (Searle, 1995), and personal and sub-personal levels (Dennett, 2010). The framework can be reconstructed in those vocabularies as well. The choice of Abhidhamma is methodological rather than privileged. It is justified by the precision of the distinction it draws and by the value for information ethics of drawing on more than one analytic tradition (Hongladarom and Bandasak, 2023).

The framework is too restrictive for future AI. Its scope is current and near-future systems, and the possibility remains open that some future system will show genuine signs of sentient volition. The immediate task is to discipline the terms applied to current systems, and the framework examines those terms rather than machine consciousness. If a future system meets the criteria for sentient moral agency, the framework would need revision, and it does not prevent that revision. It places the burden of proof where it belongs, so attributing sentient volition or moral personhood requires positive evidence rather than the force of habit. Register discipline does not replace an important account of responsibility. It comes first because if the language of attribution is left undisciplined, even a strong account will have to be applied in a confused way.

8. Implications for Documentation, Governance, and Accountability

The framework supports a more disciplined vocabulary across the technical, institutional, governance, and public practices that produce and circulate descriptions of AI systems. One principle operates across all of them. A description should report operational behaviour in functional terms, name the responsible entity, and mark any movement toward language that implies intention, remorse, or moral responsibility on the system's part. In

research documentation and peer review, the register tests can be used as a practical check on agency language because research articles help shape how AI systems are described in the scholarly literature. They can be used to ask whether claims about model behaviour use functional verbs, whether claims about responsibility name the relevant actors, and whether claims implying consciousness, intention, or moral responsibility for current systems are marked or rewritten.

The protocol can be applied directly to documentation conventions. Model cards (Mitchell et al., 2019) and datasheets for datasets (Gebru et al., 2021) can describe behaviour in operational use and locate responsibility in social-legal use. Product copy can avoid first-person responsibility claims by the system. Because the framework does not require removing personal pronouns from systems, interface design can build on it without losing the conversational style that users find comfortable. What it requires is that the deploying organisation avoid attributing remorse, responsibility, and moral judgment in the first person. In cases where acknowledging an error is needed, the organisation should frame the matter as system behaviour rather than personal apology.

The same discipline matters for what an organisation says about a system after an incident. Documentation is both a technical genre and an accountability infrastructure. A model card, an incident report, or a service description tells a later reader who acted, and that reader has nothing else to rely on. A record that places the system in subject position has therefore already performed part of the accountability judgment on the reader's behalf. Press releases, help pages, and public explanations of automated decisions carry the same burden, because they are often the only record an affected person will ever see.

In governance and law, the framework helps clarify which register a provision occupies. A provision that imposes audit obligations on a system operates in social-legal use. A provision that depends on a system deciding or intending must specify what those terms mean in context or use functional descriptions instead. The framework does not resolve disputes about legal personhood, but it clarifies the register to which they belong and warns against unmarked slippage from legal convention to moral attribution.

Across all these settings, register discipline keeps human and institutional responsibility visible in the record. It therefore addresses the moral crumple zone at the point where wording creates it (Elish, 2019), and it supports the case for drawing on more than one analytic tradition for the same problem (Hongladarom and Bandasak, 2023).

9. Conclusion

Artificial intelligence systems are commonly described as deciding, intending, rejecting, recommending, and being responsible. This article has argued that such language can be kept under discipline. It has reconstructed the distinction in Theravāda Abhidhamma between *sammuti*, or conventional designation, and *paramattha*, or ultimate moral-analytic description, as an analytic tool for describing artificial systems. The result is a Two-Register AI Agency Language Framework. Within the conventional register, the framework recognises an operational-conventional use and a social-legal-conventional use. An ultimate moral-analytic register is also recognised, one that does not justify attributing sentient volition, moral personhood, and responsibility-bearing subjectivity to current systems. Legitimating agency language within its registers, the framework blocks unmarked drift across them. Applied to three focal cases, the framework has been condensed into a language protocol that technical writers, designers, editors, regulators, and reporters can adapt.

The use of Abhidhamma is methodological. Its value lies in how clearly it separates conventional designation from morally weighty attribution, a distinction needed when a single description of a system must serve a technical purpose and an accountability purpose at once. Because the framework is limited to current and near-future systems, attributing sentient volition or moral personhood in the future will require independent argument. Since it operates on the wording through which those attributions are read, it complements rather than replaces existing accounts of responsibility gaps, moral crumple zones, role responsibility, relational standing, and legal personhood. By identifying one mechanism through which imprecise description shifts responsibility, it connects recent work on AI terminology with broader debates about moral and legal accountability.

Calling an AI system an agent is not necessarily conceptually confused. It can be a useful conventional designation, properly registered, that does not break the distinction between what a system does and what a moral person is. A register-based discipline keeps that designation available on the condition that written, published, and later relied-upon descriptions make accountability clearer rather than concealing the human and institutional responsibility behind them.

Table 1. A language protocol for AI agency language in documentation and public communication

Problematic phrase	Register confused	Why it misleads	Register-disciplined wording	Responsibility implication
"The AI decided"	Ultimate moral-analytic	Functional output is read as sentient choice	"The system selected an output under specified model, data, and deployment conditions"	Responsibility remains with developers, deployers, and oversight bodies
"The AI intended to deceive"	Ultimate moral-analytic	Functional behaviour is read as morally weighty intent	"The output produced a deceptive effect under these conditions"	Investigate design choices, training data, deployment context, and use, rather than attributing moral intent to the system
"The AI is responsible"	Social-legal and ultimate moral-analytic	Operational language slides into moral or legal liability	"The responsible parties are the humans and institutions that designed, deployed, governed, and used the system"	Allocate responsibility across designers, deployers, operators, and users
"The AI agent negotiated the contract"	Social-legal-conventional	Functional task performance is read as legal will	"The deploying organisation entered the contract through an automated negotiation system, with stated human oversight"	Liability rests with the contracting parties, not the system
"The autonomous vehicle decided to swerve"	Ultimate moral-analytic	Sentient choice is implied where there was a control output	"Under the perceived conditions, the vehicle's control system produced this manoeuvre"	Investigate manufacturer, software supplier, fleet operator, and driver responsibilities
"The AI rejected the application"	Social-legal-conventional	Moral judgment is implied where there was a score threshold	"Your application was screened out by an automated system used by [organisation], under stated criteria"	Hiring organisation is accountable, including for compliance with anti-discrimination law

Problematic phrase	Register confused	Why it misleads	Register-disciplined wording	Responsibility implication
"The AI takes responsibility for the error"	Social-legal and ultimate moral-analytic	Moral responsibility is attributed to a non-bearer	Remove the phrasing in product copy and replace it with a statement of organisational responsibility	Deploying organisation is responsible and may be liable
"I should not have said that and I apologise"	Ultimate moral-analytic	First-person remorse is staged	"This response was incorrect and has been logged for review"	Deploying organisation is responsible for content moderation and quality
"The AI understands you"	Ultimate moral-analytic	Recognition is read as comprehension	"The system produced a contextually appropriate response under these inputs"	Avoid such language in informed-consent or trust contexts
"The AI knows that this is the current guidance"	Ultimate moral-analytic	Generated text is read as testimony from a source that has checked	"The system generated this statement from the retrieved sources listed below, under stated conditions"	Provider and deploying institution are answerable for reliability and for showing the sources
"The AI found that this source is reliable"	Operational-conventional and ultimate moral-analytic	A ranking or weighting step is read as an act of assessment	"The system ranked this source highly under its stated retrieval and weighting criteria"	Provider is answerable for the criteria, and the reader retains the standing to assess the source

Note: the protocol names the register that is confused by each phrase, explains the slippage, proposes a register-disciplined alternative, and states the responsibility implication.

References

Avila Negri, S.M.C. 2021. "Robot as legal person: electronic personhood in robotics and artificial intelligence." *Frontiers in Robotics and AI* 8: article 789327. https://doi.org/10.3389/frobt.2021.789327.

Bender, E.M., T. Gebru, A. McMillan-Major, and S. Shmitchell. 2021. "On the dangers of stochastic parrots: can language models be too big?" In

Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, 610-623. New York, NY: Association for Computing Machinery. https://doi.org/10.1145/3442188.3445922.

Bodhi, B. 2007. *A Comprehensive Manual of Abhidhamma: The Abhidhammattha Sangaha of Acariya Anuruddha*. 3rd ed. Kandy: Buddhist Publication Society.

Bryson, J.J., M.E. Diamantis, and T.D. Grant. 2017. "Of, for, and by the people: the legal lacuna of synthetic persons." *Artificial Intelligence and Law* 25 (3): 273-291. https://doi.org/10.1007/s10506-017-9214-9.

Chopra, S., and L.F. White. 2011. *A Legal Theory for Autonomous Artificial Agents*. Ann Arbor, MI: University of Michigan Press. https://doi.org/10.3998/mpub.356801.

Coeckelbergh, M. 2010. "Robot rights? Towards a social-relational justification of moral consideration." *Ethics and Information Technology* 12 (3): 209-221. https://doi.org/10.1007/s10676-010-9235-5.

Coeckelbergh, M. 2014. "The moral standing of machines: towards a relational and non-Cartesian moral hermeneutics." *Philosophy & Technology* 27 (1): 61-77. https://doi.org/10.1007/s13347-013-0133-8.

Damiano, L., and P. Dumouchel. 2018. "Anthropomorphism in human–robot co-evolution." *Frontiers in Psychology* 9: article 468. https://doi.org/10.3389/fpsyg.2018.00468.

Dennett, D.C. 2010. *Content and Consciousness*. London: Routledge Classics (original work published 1969). https://doi.org/10.4324/9780203005729.

Donner, S. 2025. "What kind of documents can AI language model outputs be? The concept of artificially blended testimony." *Journal of Documentation* 81 (3): 788-809. https://doi.org/10.1108/JD-11-2024-0271.

Elish, M.C. 2019. "Moral crumple zones: cautionary tales in human-robot interaction." *Engaging Science, Technology, and Society* 5: 40-60. https://doi.org/10.17351/ests2019.260.

Floridi, L., and J.W. Sanders. 2004. "On the morality of artificial agents." *Minds and Machines* 14 (3): 349-379. https://doi.org/10.1023/B:MIND.0000035461.63578.9d.

Gebru, T., J. Morgenstern, B. Vecchione, J.W. Vaughan, H. Wallach, H. Daumé III, and K. Crawford. 2021. "Datasheets for datasets." *Communications of the ACM* 64 (12): 86-92. https://doi.org/10.1145/3458723.

Gethin, R. 1998. *The Foundations of Buddhism*. Oxford: Oxford University Press.

Gong, D. 2026. "Closing the responsibility gap: allocating responsibility according to prerequisite control and expectations for personal benefits." *Ethics and Information Technology* 28 (1): article 8. https://doi.org/10.1007/s10676-025-09883-w.

Gunkel, D.J. 2018. "The other question: can and should robots have rights?" *Ethics and Information Technology* 20 (2): 87-99. https://doi.org/10.1007/s10676-017-9442-4.

Hjørland, B. 2026. "Information retrieval or document retrieval? Terminological confusions and unrealistic goals in information science, exemplified in relation to generative artificial intelligence." *Journal of the Association for Information Science and Technology* 77 (5): 714-726. https://doi.org/10.1002/asi.70057.

Hongladarom, S. 2020. *The Ethics of AI and Robotics: A Buddhist Viewpoint*. Lanham, MD: Lexington Books. https://doi.org/10.5040/9781978735101.

Hongladarom, S., and J. Bandasak. 2023. "Non-Western AI ethics guidelines: implications for intercultural ethics of technology." *AI & Society*. https://doi.org/10.1007/s00146-023-01665-6.

Johnson, D.G. 2006. "Computer systems: moral entities but not moral agents." *Ethics and Information Technology* 8 (4): 195-204. https://doi.org/10.1007/s10676-006-9111-5.

Karunadasa, Y. 2016. *The Theravāda Abhidhamma: Its Inquiry into the Nature of Conditioned Reality*. Reprint ed. Kandy: Buddhist Publication Society.

Kozak, M. 2026. "From AI peer review to AI ethics: addressing terminological confusion in science." *Journal of Documentation* 82 (1): 111-119. https://doi.org/10.1108/JD-07-2025-0201.

Lin, C.-T. 2023. "All about the human: a Buddhist take on AI ethics." *Business Ethics, the Environment & Responsibility* 32 (4): 1257-1268. https://doi.org/10.1111/beer.12547.

Lind, M. 2025. "Alexa's agency: a corpus-based study on the linguistic attribution of humanlikeness to voice user interfaces." *AI & Society*. https://doi.org/10.1007/s00146-025-02243-8.

Matthias, A. 2004. "The responsibility gap: ascribing responsibility for the actions of learning automata." *Ethics and Information Technology* 6 (3): 175-183. https://doi.org/10.1007/s10676-004-3422-1.

Mitchell, M., S. Wu, A. Zaldivar, P. Barnes, L. Vasserman, B. Hutchinson, E. Spitzer, I.D. Raji, and T. Gebru. 2019. "Model cards for model reporting." In *Proceedings of the Conference on Fairness, Accountability,*

and Transparency, 220-229. New York, NY: Association for Computing Machinery. https://doi.org/10.1145/3287560.3287596.

Nyanaponika Thera. 1976. *Abhidhamma Studies: Researches in Buddhist Psychology*. 3rd ed. Kandy: Buddhist Publication Society.

Raptis, S.T., A.C. Kapoutsis, and E.B. Kosmatopoulos. 2025. "Agentic LLM-based robotic systems for real-world applications: a review on their agenticness and ethics." *Frontiers in Robotics and AI* 12: article 1605405. https://doi.org/10.3389/frobt.2025.1605405.

Russell, S.J., and P. Norvig. 2010. *Artificial Intelligence: A Modern Approach*. 3rd ed. Upper Saddle River, NJ: Prentice Hall.

Santoni de Sio, F., and G. Mecacci. 2021. "Four responsibility gaps with artificial intelligence: why they matter and how to address them." *Philosophy & Technology* 34 (4): 1057-1084. https://doi.org/10.1007/s13347-021-00450-x.

Scorici, G., M.D. Schultz, and P. Seele. 2022. "Anthropomorphization and beyond: conceptualizing humanwashing of AI-enabled machines." *AI & Society* 39 (2): 789-795. https://doi.org/10.1007/s00146-022-01492-1.

Searle, J.R. 1995. *The Construction of Social Reality*. New York, NY: Free Press.

Sellars, W. 1962. "Philosophy and the scientific image of man." In *Frontiers of Science and Philosophy*, edited by R. Colodny, 35-78. Pittsburgh, PA: University of Pittsburgh Press.

Senn, W.D. 2026. "The informing act: apprehension in AI-mediated systems." *Journal of Documentation 82 (6): 1646–1663*. https://doi.org/10.1108/JD-03-2026-0125.

Sharkey, A., and N. Sharkey. 2012. "Granny and the robots: ethical issues in robot care for the elderly." *Ethics and Information Technology* 14 (1): 27-40. https://doi.org/10.1007/s10676-010-9234-6.

Soulier, E.C. 2026. "Should we speak of machine agency? A case against conceptual extension." *Ethics and Information Technology* 28 (1): article 17. https://doi.org/10.1007/s10676-026-09893-2.

Sparrow, R. 2007. "Killer robots." *Journal of Applied Philosophy* 24 (1): 62-77. https://doi.org/10.1111/j.1468-5930.2007.00346.x.

Thellefsen, M.M., A. Nurma Dewi, and B. Sørensen. 2026. "The meaning of prompts: a semiotic approach to human–machine communication." *Journal of Documentation 82 (5): 1409–1429*. https://doi.org/10.1108/JD-03-2026-0140.

Tigard, D.W. 2021. "There is no techno-responsibility gap." *Philosophy & Technology* 34 (3): 589-607. https://doi.org/10.1007/s13347-020-00414-7.

Vacek, D. 2026. "Responsible artificial intelligence?" *AI & Society* 41 (2): 843-855. https://doi.org/10.1007/s00146-025-02604-3.

Wilson, P. 1983. *Second-hand Knowledge: An Inquiry into Cognitive Authority*. Westport, CT: Greenwood Press.

Wittgenstein, L. 2009. *Philosophical Investigations*. Revised 4th ed. Translated by G.E.M. Anscombe, P.M.S. Hacker, and J. Schulte. Chichester: Wiley-Blackwell.