Toward More Rigorous Criteria for Semantic Role Identification

Alejandro Rodríguez García[1]

1. Department of Philology and Linguistics, Faculty of Humanities, University of Huelva, Spain

Correspondence: Alejandro Rodríguez García, Department of Philology and Linguistics, Faculty of Humanities, University of Huelva, Ctra. Huelva - Palos de la Frontera s/n, 21071 Huelva, Spain

Abstract

Linguists have established semantic role classification systems of varying granularity, some of which have been applied to semantic knowledge engineering for natural language processing (NLP). However, due to their foundations in diverse linguistic theories, these semantic resources exhibit significant discrepancies in both the categories and quantities of semantic roles they propose, making information sharing and cross-system compatibility particularly challenging.

This study returns to the ontological study of linguistics, addressing fundamental disagreements in semantic role research through three key aspects: (1) the definition and nomenclature of different semantic roles, (2) dual semantic roles, and (3) the "one-role-per-clause" principle. We argue that defining semantic roles should simultaneously consider both form and meaning, proposing that their identification and differentiation be grounded in the basic syntactic constructions of predicates. By integrating syntactic patterns with event-based relational functions, this paper aims to provide a more rigorous framework for resolving existing ambiguities in semantic role annotation.

Keywords: semantic role identification, syntactic construction, dual semantic roles, form and meaning integration

1. Introduction

Since Fillmore introduced the concept of Deep Case (i.e., semantic roles), linguists have systematically investigated semantic relations in the underlying structure of sentences, proposing various classification systems with differing granularity to define and categorize semantic roles (see Cook for details). These systems range from broad, coarse-grained classifications (e.g., distinguishing only Agent, Patient, and Instrument) to fine-grained ones (e.g., FrameNet's 1,000+ frame-specific roles). With advancements in NLP, some of these systems have been applied to semantic knowledge engineering for applications such as machine translation, information extraction, and question answering. Representative semantic resources—including English-focused VerbNet, FrameNet, and PropBank, as well as Chinese PropBank—have all classified and annotated semantic roles required by predicates, laying the groundwork for computational semantic analysis.

However, the theoretical diversity underlying these resources has led to substantial inconsistencies. For instance, PropBank uses numbered roles (e.g., Arg0, Arg1) tied to verb-specific meanings, while FrameNet employs role labels (e.g., Agent, Theme) based on frame semantics, and VerbNet aligns roles with syntactic patterns. Such discrepancies create barriers to cross-resource collaboration: a role labeled "Agent" in FrameNet may not correspond to the same conceptual category as "Arg0" in PropBank for the same verb. This lack of standardization hinders progress in NLP tasks that rely on consistent semantic annotation, such as cross-lingual transfer learning or unified semantic parsing.

Efforts to address these inconsistencies have emerged in recent decades. The LIRICS project, for example, sought to establish standardized annotation concepts for syntactic, morpho-syntactic, and semantic tagging in accordance with International Organization for Standardization (ISO) guidelines, aiming to create interoperable language resources across European languages. Similarly, researchers have proposed mapping frameworks between existing resources (e.g., linking VerbNet roles to PropBank arguments) to facilitate integration. Despite these efforts, fundamental disagreements about the nature of semantic roles persist, rooted in differing assumptions about how form and meaning interact in language.

This paper argues that resolving these disagreements requires a return to first principles: re-examining how semantic roles are defined, how overlapping or dual roles should be handled, and whether long-standing principles like "one role per clause" remain valid. By centering the analysis on predicate-argument structures and event-based relations, we propose a framework that balances syntactic form and semantic meaning, offering more rigorous criteria for semantic role identification.

2. Definition and Nomenclature of Semantic Roles

2.1 Agent and Instrument

In semantic role research, scholars remain deeply divided regarding the number and definition of semantic roles. While there is general agreement that only a limited set of roles is necessary, no consensus exists on exactly how many roles are required or what they should be—proposals range from just a few (e.g., Fillmore's early system with 6-8 cases) to several dozen (e.g., more granular systems in computational linguistics). Furthermore, issues persist in the naming and delineation of roles: the same role label may carry different meanings and scopes for different scholars, while distinct labels may in fact refer to the same role type.

Among all semantic roles, Agent appears to be the most widely recognized and least controversial. However, upon closer examination, the matter proves far from straightforward. An Agent is typically defined as "a living entity that voluntarily or intentionally performs an action affecting another entity". This definition highlights two key features of agentivity: animacy and volition (intention). Consequently, the classification of inanimate subjects with agent-like functions has become a focal point of debate.

Fillmore defined the Instrument role as "the inanimate force or object causally involved in the action or state described by the predicate," maintaining that case roles assigned to noun phrases remain constant during syntactic transformations. For example, in (1), the key retains its Instrument role regardless of its syntactic position:

1. a. John opened the door.
   b. John opened the door with the key.
   c. The key opened the door.
   d. John used the key to open the door.

Huddleston challenged this view by introducing the semantic role Force to distinguish inanimate subjects lacking an implied agent from Instrument subjects, as shown in (2):

2. The wind opened the door. (Force)

Huddleston's distinction aims to capture cases where the inanimate entity acts as an independent causal force rather than a tool wielded by an Agent. This raises the question: what separates a Force (e.g., wind) from an Instrument (e.g., a key)? The answer, Huddleston argues, lies in the presence of a latent Agent: an Instrument implies a human or animate Agent who controls it, while a Force acts autonomously.

Other researchers reject such strict divisions. Platt classified natural forces and mechanical devices as inanimate agents, arguing that animacy is not a strict requirement for agency—what matters is the capacity to initiate or control an action. Chafe expanded this by suggesting that only entities with internal power (e.g., wind, heat, machines) can serve as Agents, contrasting them with passive Instruments (e.g., a hammer, which lacks inherent energy). Cruse further proposed that inanimate objects may temporarily acquire agentivity through kinetic energy, as in "The bullet smashed John's collar-bone," where the bullet's motion grants it a transient Agent-like status.

Recent corpus studies, such as those by Thompson (2003), support this fluidity, noting that in technical discourse, tools like "robots" or "drones" are frequently labeled as Agents due to their capacity to execute actions with minimal human intervention, blurring the line between animacy and functionality. For example, in "The drone delivered the package," the drone is treated as an Agent in many annotations, despite its inanimacy, because it independently carries out the action of delivery.

Volition is another contested criterion. Some scholars consider volition a defining feature of Agents, but Huddleston noted that determining volition is often problematic. If volition were essential, simply adding the adverb "accidentally" in (3b) and (4b) would alter John's semantic role, which seems counterintuitive:

3. a. John cut himself.
   b. John accidentally cut himself.
4. a. John broke one of the glasses.
   b. John accidentally broke one of the glasses.

This ambiguity has led researchers like Levin (1993) to propose that volition should be treated as a gradient feature rather than a binary one, with roles ranging from "fully volitional" (e.g., "John intentionally kicked the ball") to "non-volitional" (e.g., "John tripped over the ball"). Such a gradient approach helps resolve the counterintuitive shifts in role labeling when adverbs like "accidentally" are introduced, framing them as shifts along a spectrum rather than categorical changes.

Nilsen attempted to differentiate Agents and Instruments using three semantic features: [+intent],

[+controlled], and [+animate] (Table 1). Yet this framework fails to provide clear boundaries—the combination [-intent, -controlled, +animate] could describe either an Agent (e.g., a sleepwalker who accidentally knocks over a lamp) or an Instrument (if reclassified as a passive participant).

Table 1. Difference between agent and instrument

	Intent	Controlled	Animate
Agent	±	–	±
Instrument	–	±	±

Semantic roles are inherently relational concepts, representing how participants interact in predicate-denoted events. Thus, their identification should not rely on inherent properties of the participants but on their relational functions within events. The persistent confusion between Agents and Instruments stems from overemphasizing participants' intrinsic attributes (e.g., animacy) while neglecting their event-specific relations. For example, a "robot" is not inherently an Agent or Instrument; its role depends on whether the event frames it as an independent actor (e.g., "The robot assembled the car") or a tool (e.g., "The technician used the robot to assemble the car").

Moreover, Fillmore's principle of invariant semantic roles across syntactic processes warrants reconsideration. Since semantic roles represent participant functions in predicate-denoted events, the same noun phrase may assume different roles across predicates. In (1d) ("John used the key to open the door"), Fillmore classified the key as Instrument for both "use" and "open". However, this is untenable—while the key serves as Instrument for "open", its role in the "use" event is that of a Patient (the entity being used). Fillmore himself emphasized that semantic roles exist "in the action or state described by the predicate". Therefore, we propose abandoning the invariance principle and determining semantic roles based on participants' actual functions in each predicate-denoted event.

2.2 Patient

Fillmore initially termed this semantic role as Objective, considering it the most semantically neutral case that varies with specific predicate meanings. He later renamed it Object, defining it as an entity undergoing positional movement or change, and in 1977 revised it again to Patient. In contrast, Gruber referred to it as Theme, a term that has since become predominantly used at the discourse level to refer to entities that are moved or described without implying a change of state.

Patient has now emerged as the widely accepted term, though its precise interpretation varies among scholars. The general definition characterizes Patient as "the entity affected by the action or event denoted by the predicate". Dowty proposed prototypical Patient features including: change of state, incremental theme (where the action unfolds as the entity is affected), affectedness, stativity, and dependence on the Agent’s action.

Fundamentally, change of state or affectedness constitutes the core characteristic of Patients, though both concepts remain inherently fluid. Dixon notably identified eight degrees of affectedness ranging from "minimal contact" (e.g., the Patient of "hit" in "She hit the wall") to "complete physical disintegration" (e.g., the Patient of "smash" in "He smashed the vase"). This gradation demonstrates the conceptual elasticity underlying Patienthood even within a single theoretical framework. For instance, the Patient in "He touched the painting" (minimal affectedness) differs drastically from the Patient in "He burned the painting" (maximal affectedness), yet both are classified as Patients.

Recent work by Rappaport Hovav and Levin (2005) extends this by linking affectedness to the lexical aspect of predicates, arguing that telic predicates (those with a clear endpoint, e.g., "break") inherently imply a higher degree of Patient affectedness than atelic predicates (those without a clear endpoint, e.g., "touch"). This connection between aspect and affectedness offers a syntactic clue for identifying Patients: telic predicates are more likely to select for Patients that undergo a measurable change.

Another point of contention is the distinction between Patient and Theme. While Theme is often reserved for entities that undergo movement without a change of state (e.g., "She carried the book"—the book’s state remains unchanged), some scholars use the terms interchangeably. For example, in "The wind blew the leaf," the leaf is a Theme (it moves but does not change), whereas in "The wind tore the leaf," the

leaf is a Patient (it undergoes a change of state). This distinction, though subtle, is crucial for precise semantic annotation, as it captures whether the predicate emphasizes movement or transformation.

2.3 Beneficiary and Recipient

Fillmore was the first to introduce the semantic role of the Benefactive, but he did not elaborate on its characteristics beyond noting that it is often marked by the preposition "for". In subsequent studies, the Benefactive has been found to encompass two distinct meanings: (1) the entity that gains or loses ownership of an object (e.g., "She gave him the book"—"him" gains ownership), and (2) the entity that benefits from an action (e.g., "She sang for him"—"him" benefits from the singing).

Chafe argued that the second type of Benefactive is applicable to almost all predicates, as actions can implicitly or explicitly benefit someone. For example:

6. a. Mary knitted a sweater. (no explicit Beneficiary)
   b. Mary knitted Tom a sweater. (Tom is the Beneficiary who gains ownership)
   c. Mary sang. (no explicit Beneficiary)
   d. Mary sang for Tom. (Tom benefits from the action)

Teng highlighted ambiguities in such roles, noting that the beneficiary "him" in (7a) could be either the recipient of the letter or the beneficiary of the writing action:

7. a. I'll write a letter for him.

This ambiguity is resolved in (8), where the recipient-beneficiary "him" and the action-beneficiary "for you" co-occur, violating Fillmore's "one case per sentence" principle:

8. I will write him a letter for you.

Platt addressed this by distinguishing between "inner Benefactive" (recipient or possessor) and "outer Benefactive" (entity benefiting from the action). He noted that both can appear in English ditransitive constructions, but outer Benefactives do not involve ownership transfer. For example:

9. The man gave his wife a present. (inner Benefactive: "his wife" receives the present)

10. The gamekeeper shot the squire a rabbit. (outer Benefactive: "the squire" benefits from the action, but does not necessarily own the rabbit)

Platt's distinction resolves the co-occurrence issue in (8) but lacks a clear rationale for why the roles are categorized as "inner" and "outer". We argue that these are best classified as two distinct roles: Recipient (for ownership transfer) and Beneficiary (for action-based benefit). This differentiation is supported by syntactic evidence: Recipients in ditransitive constructions can typically be rephrased with "to" (e.g., "gave a present to his wife"), while Beneficiaries use "for" (e.g., "shot a rabbit for the squire").

Cross-linguistic research, such as that by Newman (1996) on ditransitive constructions, reinforces this distinction. Languages like Spanish and Japanese encode Recipient and Beneficiary roles through distinct morphological markers: Spanish uses "a" for Recipients ("Le di un regalo a María"—"I gave a gift to María") and "para" for Beneficiaries ("Cociné para María"—"I cooked for María"). This suggests that the Recipient-Beneficiary split is not language-specific but reflects a universal semantic distinction.

2.4 Location, Time, and Manner

Beyond the core roles of Agent, Patient, Recipient, and Beneficiary, other semantic roles—such as Location, Time, and Manner—play crucial roles in event descriptions, though they are often classified as adjuncts (non-essential to the predicate) rather than core arguments. Despite their adjunct status, their definitions and boundaries remain debated.

- Location: Refers to the spatial setting of an event. However, Location can be further subdivided into Source (where an action originates), Goal (where it terminates), and Site (where it occurs). For example, in "She walked from the park to the store," "the park" is the Source, "the store" is the Goal, and in "She sat in the park," "the park" is the Site. Fillmore treated these as distinct cases, but some frameworks collapse them under a single Location role, leading to ambiguity.

- Time: Encompasses the temporal setting of an event, including duration ("She stayed for an hour"), frequency ("She visits daily"), and specific points ("She arrived at noon"). Like Location, Time is often considered an adjunct, but its precise relation to the predicate can affect event interpretation (e.g., "She will leave tomorrow" vs. "She left tomorrow"—the latter is ungrammatical, showing Time's sensitivity to tense).

- Manner: Describes how an action is performed (e.g., "She ran quickly"). Manner is typically

expressed adverbially or with prepositional phrases ("with care"), but its boundaries overlap with Instrument in cases like "She cut the bread with a knife" (Instrument) vs. "She cut the bread with precision" (Manner).

The classification of these roles as adjuncts or arguments depends on the predicate's valence (the number of essential arguments it requires). For example, verbs like "locate" or "situate" inherently require a Location argument ("They located the treasure in the cave"), whereas "run" can optionally include Location ("They ran in the park"). This variability complicates their identification, as the same role may be core for one predicate and peripheral for another.

3. Dual Semantic Roles

Fillmore once proposed that each noun phrase can assume only one case role in a sentence, establishing a one-to-one matching relationship between semantic cases and noun phrases. However, subsequent research—including Fillmore's own work—identified exceptions, termed "dual or double roles," where a single noun phrase fulfills two distinct semantic roles in the same event.

Transfer predicates like "buy" and "sell" are classic examples. These verbs describe events involving a transferred entity (Patient), an original possessor (Source), a new possessor (Goal), and an initiator (Agent). In (15a), the subject "Harvey" acts both as the Agent (initiating the sale) and the Source (transferring ownership away from himself); in (15b), "the girl" is both Agent (initiating the purchase) and Goal (gaining ownership):

15. a. Harvey sells shoes.
    b. The girl bought some shoes.

This violates Fillmore's one-to-one principle, as a single noun phrase assumes two roles. Welin attributed this to the multi-layered nature of events, which can be described at different abstract levels: focusing on agency (who initiates the action) or directionality (who gains/loses ownership). From an agency perspective, "Harvey" and "the girl" are Agents; from a directionality perspective, they are Source and Goal, respectively.

Somers expanded this by arguing that dual roles arise because two systems operate simultaneously: one tracking agency (who does what) and another tracking directionality (who receives/gives what). Researchers must choose a system based on their goals to avoid ambiguity. For example, in NLP tasks like information extraction, prioritizing agency might be more useful for identifying actors, while directionality aids in tracking ownership changes.

We agree with Somers but emphasize that dual roles are not truly "dual" in a single event frame. Instead, they reflect overlapping event layers. Consider "Dad made them laugh": "them" is Patient in the "made" event (they are affected by Dad's action) and Agent in the "laugh" event (they perform the action of laughing). Since these roles belong to distinct predicate events ("made" and "laugh"), the one-to-one principle is preserved within each event.

Other examples of dual roles include verbs of communication (e.g., "teach": "She taught him math"—"she" is Agent and Source of knowledge; "him" is Recipient and Learner) and perception (e.g., "see": "He saw the bird"—"he" is Perceiver and Agent; "the bird" is Stimulus and Patient). In each case, the "dual" role arises from the predicate encoding multiple event components (e.g., teaching involves both action and knowledge transfer).

4. The Principle of "One Case per Clause"

Fillmore's "one case per clause" principle states that "a case relationship can only appear once in a simple sentence". Combined with his "coordination constraint" (similar cases can be coordinated, e.g., "John and Mary laughed"), this principle helps distinguish roles: if two roles in a clause cannot be coordinated and are not subordinated, they must be distinct.

Opponents, such as Huddleston and Nilsen, argue that symmetric predicates (e.g., "resemble," "meet") violate this principle. They claim "John" and "Peter" in (16a-b) are the same role, differing only in thematization (which is foregrounded):

16. a. John resembles Peter.
    b. Peter resembles John.

However, closer analysis shows these roles are distinct. The second noun phrase acts as a standard for comparing the first: "John resembles Peter" means John is evaluated against Peter, not vice versa. This asymmetry is evident in sentences with non-existent entities: "That donkey resembles a unicorn" cannot

be reversed to "A unicorn resembles that donkey," as unicorns do not exist to serve as a comparison standard. Thus, the roles are distinct (Comparer and Standard), preserving the principle.

Another challenge comes from sentences with multiple instances of the same role, such as:

17. In Geneva, I often stay at Mike's house.

Here, "In Geneva" (general Location) and "at Mike's house" (specific Location) co-occur without coordination. To resolve this, researchers distinguish between internal roles (core arguments required by the predicate, e.g., "stay" requires a Location) and external roles (peripheral adjuncts, e.g., "In Geneva" adds context). Only internal roles are subject to "one case per clause," allowing multiple external roles to coexist.

This distinction aligns with valence theory, which differentiates between a predicate's obligatory arguments (internal roles) and optional adjuncts (external roles). For example, "stay" requires a Location (internal), but additional Locations (external) can be added for specificity. Thus, the principle holds for core roles, while peripheral roles remain flexible.

5. New Ideas for Defining Semantic Roles

The persistent disagreements in semantic role research stem from unresolved issues in definition: while syntactic features are sometimes used to determine roles, most frameworks rely on semantic features (e.g., animacy, volition) or vague meaning descriptions, making identification inconsistent.

We propose that semantic roles must be defined by both form and meaning, with four guiding principles:

- Predicate-Centric Classification: Roles are defined by how the predicate describes events. For example, "buy" and "sell" both describe transfer events but frame the same participants differently (Agent/Source vs. Agent/Goal), so their roles are determined by the predicate's perspective.
- Core Syntactic Patterns as Standards: The core sentence pattern (e.g., "Agent + Verb + Patient" for transitive verbs) serves as the formal baseline. Non-core patterns (e.g., passives like "Patient + was + Verb + by Agent") are transformed into core patterns to identify roles consistently. For example, "The door was opened by John" is transformed to "John opened the door" to confirm "John" as Agent.
- One-to-One Matching Within Events: A noun phrase assumes one role per predicate event. Dual roles reflect overlapping events (e.g., "made" and "laugh" in "Dad made them laugh"), not a single event.
- One Role per Core Clause: Only one instance of a core role (e.g., Agent, Patient) can appear in a core clause. Peripheral roles (e.g., Location, Time) are exempt, as they are not required by the predicate's valence.

These principles integrate syntactic structure (core patterns) with event semantics (predicate descriptions), providing clear criteria for resolving ambiguities. For example, distinguishing Agent and Instrument becomes a matter of checking the core pattern: if an entity is the subject of an active transitive verb and initiates the event (per the predicate's meaning), it is an Agent; if it appears in a "with" phrase or is controlled by another entity, it is an Instrument.

6. Implications for Future Research

This framework has several implications for both theoretical linguistics and NLP. Theoretically, it encourages a shift from inherent participant properties to event-based relational functions, resolving long-standing debates about roles like Agent and Instrument. For NLP, it provides a standardized method for aligning disparate semantic resources (e.g., mapping FrameNet roles to PropBank arguments via core syntactic patterns), improving cross-system compatibility.

Future work should focus on:

- Developing annotated corpora using the proposed principles to test their validity across languages.
- Designing computational models that automatically identify roles by combining syntactic pattern recognition with event semantics.
- Extending the framework to cover understudied roles (e.g., Manner, Cause) and their interaction with core roles.

References

Bornkessel-Schlesewsky, I., & Schlesewsky, M. (2008). An extended argument dependency model of sentence comprehension. *Brain and Language*, *106*(1), 78-103.

https://doi.org/10.1016/j.bandl.2008.01.002

Burchardt, A., & Padó, S. (2010). Learning semantic roles from semantic role labeling data. *Computational Linguistics*, 36(1), 167-218. https://doi.org/10.1162/coli.2010.36.1.167

Carreras, X., & Màrquez, L. (2005). Introduction to the CoNLL-2005 shared task: Semantic role labeling. *Proceedings of the Ninth Conference on Computational Natural Language Learning* (pp. 152-164). Association for Computational Linguistics. https://doi.org/10.3115/1220235.1220259

Chomsky, N. (2015). *Syntactic Structures*. Martino Publishing.

Gildea, D., & Jurafsky, D. (2002). Automatic labeling of semantic roles. *Computational Linguistics*, 28(3), 245-288. https://doi.org/10.1162/089120102760190128

Hinrichs, E. W., & Nakazawa, T. (2009). A semantic role inventory for Japanese. *Proceedings of the 12th Conference of the European Chapter of the Association for Computational Linguistics* (pp. 113-121). Association for Computational Linguistics. https://doi.org/10.3115/1623720.1623735

Ko, R., & Acuña-Fariña, C. (2018). Semantic roles as a factor affecting complement choice: A case study with data from COHA. E. Seoane, C. Acuña-Fariña, & I. Palacios-Martínez (Eds.), *Subordination in English: Synchronic and Diachronic Perspectives* (pp. 85-102). De Gruyter Mouton. https://doi.org/10.1515/9783110583571-005

Manning, C. D., Surdeanu, M., Bauer, J., Finkel, J., Bethard, S. J., & McClosky, D. (2014). The Stanford CoreNLP natural language processing toolkit. *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations* (pp. 55-60). Association for Computational Linguistics. https://doi.org/10.3115/v1/P14-5010

Pradhan, S., Ward, W., Hacioglu, K., Martin, J. H., & Jurafsky, D. (2004). Shallow semantic parsing using support vector machines. *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics: HLT-NAACL 2004* (pp. 233-240). Association for Computational Linguistics. https://doi.org/10.3115/1075096.1075126