Title

Perceptually Grounding and Determining Truth Values of Propositions Joined by Logical Connectives: A Neural Process Account

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 48(0)

Authors

Akinci, Leyla Betül
Kang, Minseok
Schöner, Gregor

Publication Date

2026

Perceptually Grounding and Determining Truth Values of Propositions Joined by Logical Connectives: A Neural Process Account

Leyla B. Akinci (lb.akinci@edu.rptu.de)[1], Minseok Kang (minseok.kang@ini.rub.de)[2] & Gregor Schöner (gregor.schoener@ini.rub.de)[2]

[1]Center for Cognitive Science, RPTU, Germany

[2]Institute for Neural Computation, Faculty of Computer Science, Ruhr-University Bochum, Germany

Abstract

When descriptive sentences are linked by disjunctive logical connectives (or, xor), they may be true for multiple possible world configurations, unlike sentences linked by the conjunctive which reference a single possibility. Moreover, some of the component propositions may be false in some of the possible world configurations, while the compound sentence is still true. Determining the truth values of such sentences requires, therefore, dealing with propositions that cannot be grounded because they are false and evaluating sentences relative to multiple possible outcomes of the grounding process. Building on Dynamic Field Theory (DFT), we propose a neural process account that (1) represents affirmative and negative relational propositions linked by logical connectives, (2) autonomously drives perceptual grounding via sequential attention shifts and evaluations, and (3) determines the truth values of each component proposition and of the compound sentence by matching grounding outcomes to the possibility representations implied by each connective. The model reproduces the negation-by-truth-value interaction effect at the proposition level and extends it to the compound level.

Keywords: dynamic field theory; grounding; neural dynamics; logical connectives; computational modelling; process model; connectionist model;

Introduction

When someone states that the keys are “on the table (and) to the right of the vase”, a listener may be able to find them and thus perceptually ground this nested relational sentence. The truth value of the statement is judged based on whether they found the keys or not, and it requires simply testing one hypothesis that entails the possibility of the keys being in one particular location defined relatively to the location of two other objects. If, instead, the listener is told that the keys are “either on the table or in the bag”, he is led to consider two possibilities, only one of which may be true. Perceptually grounding these types of sentences requires handling and representing different possibilities in an abstract way.

Understanding how sentences that contain logical connectives such as *or*, *either-or* (xor), or *and* are processed by humans may advance the program of grounded and embodied cognition Barsalou, 2008) which postulates that all cognition, including higher cognition, ultimately emerges from sensory-motor experience. This postulate is central to Dynamic Field Theory (DFT; Schöner et al., 2016) in which cognitive processes are generated as time courses of neural activation states that interact with sensory input and drive motor output.

Previously, Richter et al. (2021) accounted for the perceptual grounding of spatial relations within that framework. This was extended by Sabinasz and Schöner (2023) to the perceptual grounding of nested relational phrases conjoined by *and*. Recently, Kati et al. (2024) demonstrated the perceptual grounding of propositions of affirmative or negative polarity and proposed a framework for determining truth values within such process models.

The current work aims to extend this line of modeling to ground the semantics of logical connectives *or*, *xor*, and *and* in tasks that require truth value evaluation based on visual input. We propose a neural process model that represents the different possible outcomes of attempts to perceptually ground spatial propositions. On this basis, the model evaluates the truth value of compound sentences that link spatial propositions by logical connectives.

Methods

Dynamic Field Theory (Schöner et al., 2016, DFT) is a mathematical framework to understand cognition in terms of the underlying neural processes. In DFT, cognitive states are instantiated as distributions of activation in neural populations that evolve continuously in time under the influence of recurrent connectivity (“interaction”) and coupling to sensory and motor systems. These distributions are modeled as time, t, dependent activation fields, $u(x,t)$, over low-dimensional continuous feature spaces, x (such as the position and the hue of localized patches on visual space), to which a field may be tuned by virtue of its feed-forward connectivity from the sensory surface. The time courses of activation fields, $u(x,t)$, emerge from neural dynamics of this generic form:

$$\tau\,\dot{u}(x,t) = -u(x,t) + h + s(x,t) + w_\xi\,\xi(x,t) + \int k(x-x')\,g(u(x',t))\,dx'. \tag{1}$$

Activation states, $u(x,t)$, of the fields evolve on the time scale, $\tau = 100$ ms in the present model, integrates external input $s(x,t)$ with recurrent connectivity described by the kernel $k(x-x')$, a weighted sum of local excitatory interaction, lateral and global inhibitory interaction, driven by activation from anywhere in the field, x', and passed through a sigmoidal function, $g(u(x',t))$ whose threshold is zero. In the absence of external input, the resting level, $h < 0$, of the field is a stable fixed point. Variability arises from Gaussian, white stochastic inputs, $\xi(x,t)$, whose variance is set by w_ξ. Fields

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

may be defined over feature spaces, x, ranging in dimensionality from 1 to 4. Zero-dimensional fields, also called neural nodes, encode discrete concepts such as "true" or "red". DFT architectures consist of multiple coupled fields and nodes. Sources of coupling are the outputs, $g(u(x,t))$, of given fields that provide input to the target fields. The dimensionality of the source of a connection may be expanded or contracted to match that of the target of the connection.

The sub-threshold attractor, $s(x,t)+h<0$, becomes unstable in the detection instability when pushed through zero by increasing external input. Localized peaks of activation are supra-threshold attractors to which the field then converges. Depending on conditions and parameters, a single peak may emerge (selection), or multiple peaks may form the attractor. Peaks may persist when input is removed (sustained).

Stimuli and Task

The model is given compound sentences composed of relational propositions linked by logical connectives. Its task is to evaluate the truth value of the compound sentence based on the truth values of individual propositions determined on the basis of the given visual scene.

Visual and language inputs to the model are illustrated in Fig. 1. Visual stimuli consisted of a white background with three filled colored circles (blue, red, and green) in varying spatial configurations that matched different spatial relational descriptions (above, below, left of, right of). Language stimuli contained two propositions about spatial relations that were linked by varying logical connectives (OR, XOR, AND) and had two possible polarities (positive, negative). Together, this led to 18 conditions, with each logical connective in 6 conditions given by combinations of polarity and truth-values of individual propositions (TT, TF/FT, FF).

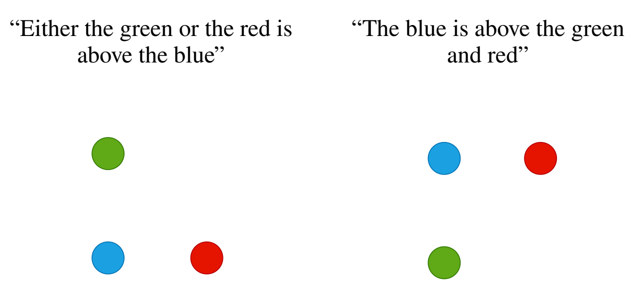

Figure 1: Two examples of visual-language stimuli: All propositions have positive polarity, one is true, the other false. The compound sentence is true for the left and false for the right.

The DFT Architecture

The architecture (Figure 2) is one large dynamical system that consists of coupled integro-differential equations of the type of Eq. 1. We break down the architecture into five main functional sub-architectures. The *neural representation of propositions* (orange areas) consists of neural nodes representing concepts activated by language input, which is sustained as a working memory of each proposition. Propositions contain color concepts in two roles, reference and target, to signify objects and spatial relations in two polarities, affirmative or negated. Multiple copies of these nodes exist to represent different propositions (see Sabinasz and Schöner, 2023), here restricted to only two: P1 and P2 (grey sheets staggered in depth). The logical connectives (or, xor, and) are symmetrical and relate these two propositions.

The *perceptual grounding* (green areas) detects whether objects referred to by language exist in the visual scene (cf. Richter et al. (2021) for details). The model first performs visual search for objects cued by the target and reference colors. Their spatial locations are represented by sustained peaks in the target and reference fields. A coordinate transformation centers the target on the location of the reference object in the Relational CoS field, to which a synaptic pattern is applied that encodes the activated spatial relation. A peak forms if the centered target falls into the excitatory region of the synaptic relational pattern. This activates the CoS node signaling successful grounding of the relation. CoS stands for *condition of satisfaction* as explained below. Failure to form a peak signifies failure to ground the relation. To detect this non-event, the *Condition of Dissatisfaction* (CoD) node is activated by default: It evolves on a slower time scale ($\tau = 130ms$) so that it is activated unless inhibited by the activation of the CoS node.

The *proposition truth evaluation* sub-architecture (red areas), evaluates the truth values of the propositions based on the outcome of their perceptual grounding as signaled by CoS or CoD nodes (cf. Kati et al. (2024) for a full account). For an affirmative proposition, CoS signals truth; for a negative proposition, CoD signals truth. This sub-architecture represents all possible combinations of grounding result and polarity through four nodes: True Affirmative ($O+$), False Affirmative ($X+$), False Negative ($O-$), and True Negative ($X-$). Two nodes represent the truth value of a proposition, true T, or false F, with each proposition having its own mental model (grey sheets into which the red areas are embedded).

The *sentence truth evaluation* sub-architecture (purple areas) consists of five neural nodes. Three of these represent the possibilities entailed by the logical connectives that link a pair of propositions (TF, FT, and TT) (FF is not entailed by the connectives we consider, as it holds false for all of them). The remaining two nodes represent the truth value of the compound sentence. The truth value of the sentence is determined by detecting a match between the neural representation of the logical connectives, given as the possibilities they entail, and the truth values of the individual propositions. Detecting a match boosts the True node, whereas the absence of a match drives the False node, analogously to how CoS and CoD signal success or failure in perceptual grounding.

Finally, the autonomous generation of *sequences of pro-*

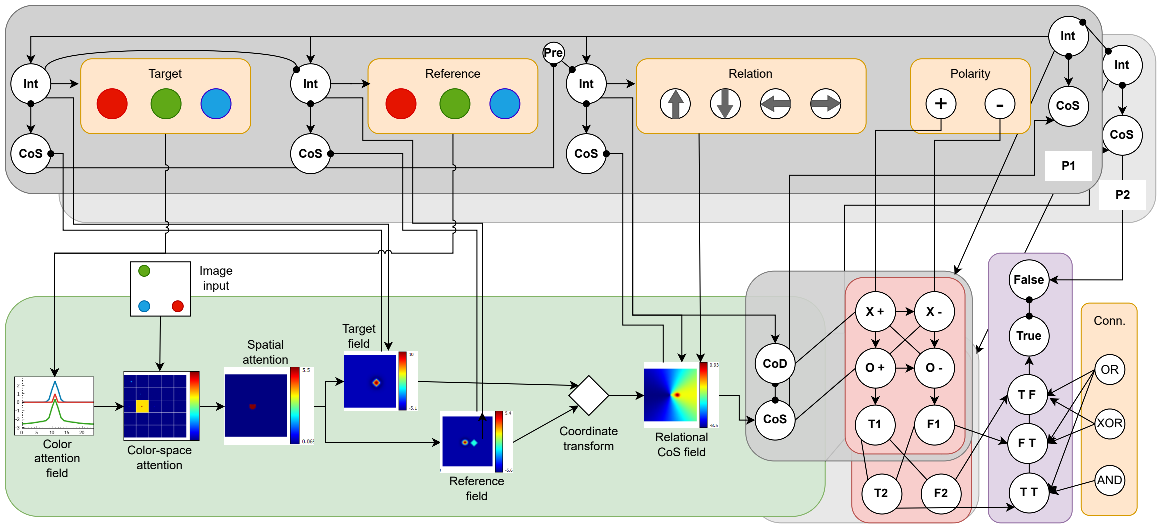

Figure 2: A schematic of the neural dynamic architecture. Background colors highlight sub-architectures described below. Connections marked by pointed arrowheads are excitatory, connections marked by rounded arrowheads are inhibitory.

cessing steps emerges from coupled Intention and Condition of Satisfaction (CoS) nodes at two levels. Within each proposition, they organize the sequential grounding of target object, reference object, and relation (top of Figure 2). The sequential grounding of different propositions is organized by the coupling among such pairs (top right of Figure 2). In each case, an Intention node activates a particular process by boosting the relevant neural fields until its successful completion is signaled by the associated CoS node, which in turn deactivates the Intention node. The coupling among intention nodes determines which other Intention node may next be activated.

The termination of a perceptual grounding act, successful or not, is detected by CoS and CoD nodes (bottom right in Figure 2) as explained earlier. These induce activation of appropriate interneurons in the sentence truth evaluation subarchitecture, transforming success or failure into truth values of the current proposition.

Results

The model is implemented for numerical simulation in CEDAR (Lomp et al., 2016). We first demonstrate how activation time courses in relevant neural fields autonomously generate the relevant grounding and truth determination processes. We then use reaction times obtained from the model to characterize different categories of tasks and make predictions that extend the known negation-by-truth-value interaction effect. The model implementation, reaction time data and its processing, visual and language stimuli can be found in the corresponding repository[1].

The time courses of activation shown in Figure 3 arise during the processing and perceptual grounding of the sentence "either the green or the red is above the blue" while the model is provided with the visual stimulus shown on the left of Figure 1. Proposition P1 ("green above blue") is true, proposition P2 ("red above blue") is false. The compound sentence, P1 *xor* P2, is true.

At the outset, the P1 intention node (blue bar in Fig. 3) is above threshold and activates a moment later the intention node for grounding the first target object (violet bar). At t_1, an object matching the current target color *green* has been selected, represented by a sustained activation peak in the target field. This peak activates the target CoS node (pink bar), which inhibits the target intention node (violet bar), terminating target grounding. A moment later, the intention node for grounding the first reference object (grey bar) is activated.

An analogous process then selects a reference object, creating an activation peak in the reference field, t_2. The activated CoS nodes of target (pink bar) and reference (cyan bar) grounding inhibit the precondition node (green bar), which releases the relation intention node (red bar) from inhibition. Very quickly, the now boosted relational CoS field forms a peak as input from the target field, centered on the reference object through a coordinate transform, overlaps with the region that the *above* relation node excites. All this happens quickly at t_3, leading to an activated P1 relational CoS node (blue bar) that signals the success of grounding while the associated CoD node is prevented from activating.

This ends the intention to ground proposition 1 (via some intermediate steps) and releases the intention to ground proposition 2 (orange bar) from inhibition at t_4. The grounding of target 2 (brown bar) succeeds at t_5, the grounding of reference object 2 (yellow bar) succeeds at t_6. At t_7, the relation intention node (red bar) is thus activated for a second time, boosting the relational CoS field and the CoD node. The second proposition is false, however, so no peak forms in the relational CoS field as the target input does not overlap with the region that

[1] https://github.com/LeylaBA/grounding_logical_connectives.git

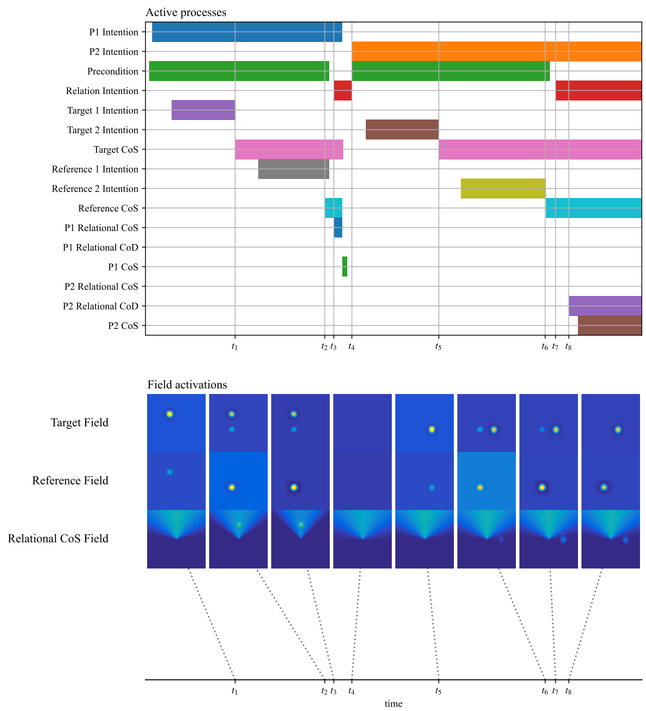

Figure 3: Time course of grounding two propositions. The upper panel indicates by solid bars time intervals during which neural nodes designated on the left have above-threshold activation. The transitions from sub- to supra-threshold activation occurring through detection instabilities are marked as event times, t_i, for some nodes. Activation patterns in relevant fields at these event times are shown in the lower panel. Blue indicates below threshold activation, orange-yellow above threshold activation.

the *above* node excites. Failure to activate the relational CoS node allows the CoD node, which operates on a slower time scale, to reach activation at t_8, signaling the failure to ground the proposition. Therefore, processing proposition two took longer than processing proposition one (orange vs blue bar at the top of Fig. 3).

As the perceptual grounding proceeds, the truth values of the propositions are assessed in parallel, leading seamlessly to the evaluation of the compound sentence. Figure 4 illustrates the time course of some nodes in the sentence truth evaluation system from that same simulation. The *xor* node (orange bar) is active from the start. The + nodes in the proposition truth evaluation (compare Fig. 2) that represent truth values of affirmative phrases are boosted because the first proposition is affirmative. As the processing of proposition 1 succeeds at time t_3, its CoS node activates the $O+$ node (O for true, X for false). This leads to activation of T1 (red bar in Fig. 4), signaling that the first proposition has been determined as true.

Similarly, for the second proposition, also affirmative, the + nodes are boosted. This time, failure to ground the relation leads to the CoD node becoming active at t_8, which activates

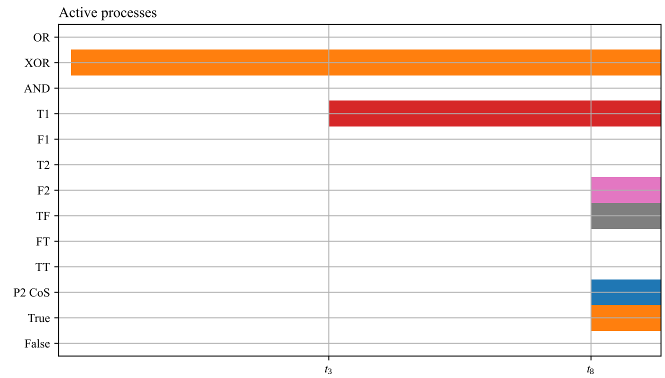

Figure 4: Time courses of some nodes involved in the truth evaluation of a sentence conjoined by *xor*. Colored bars indicate times during which nodes are above threshold.

the $X+$ node and then the F2 node (pink bar in Fig. 4), signaling that the second proposition has been determined as false.

Thereafter, the truth value of the compound sentence is determined. The *xor* node has boosted the TF and FT nodes. At the end of grounding signaled by the P2 CoS node (blue bar in Figure 4), the active T1 and F2 nodes bring the TF node (gray bar) above threshold. The node representing that the compound sentence is true (orange bar at the bottom) is activated, which stops the node representing that the compound sentence is false from activating by default.

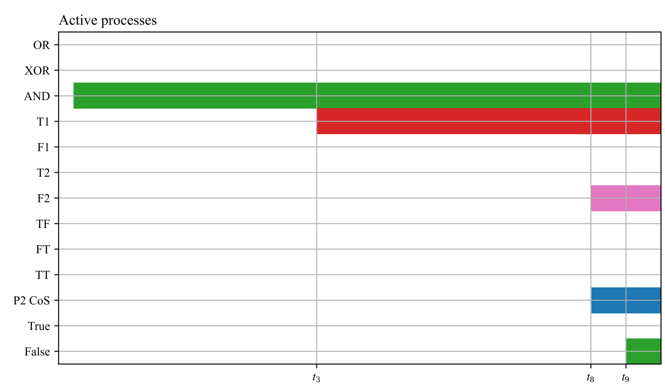

Figure 5: Time courses of some nodes involved in the truth evaluation of a sentence conjoined by *and*.

To demonstrate a false affirmative case on the sentence level, Figure 5 depicts the temporal progression of the grounding of propositions from the stimulus pair on the right in Figure 1. This stimulus pair also yields TF truth values for propositions, but this time F for the overall sentence, as it involves *and*. Here, the same possibility node TF does not receive activation from the AND concept node (green) because it does not entail it as possible. Hence, even when T1 and F2 gradually become active and project onto TF, it doesn't suffice to activate this node (or any of the possibility nodes for that matter). Once the F2 becomes active, the delayed activation of the false node (second green), initiated by the P2 CoS node (blue), prevails, indicating none of the possible worlds match the real world.

ID	LC	pol	prop	sent	nCoD	Mean	SD
1	OR	+	TT	T	0	11960	612
2	OR	+	TF	T	1	12360	382
3	OR	+	FF	F	3	13350	449
4	OR	−	TT	T	2	12750	528
5	OR	−	TF	T	1	12370	541
6	OR	−	FF	F	1	12320	406
7	XOR	+	TT	F	1	12650	486
8	XOR	+	TF	T	1	12280	411
9	XOR	+	FF	F	3	13340	400
10	XOR	−	TT	F	3	13270	501
11	XOR	−	TF	T	1	12230	428
12	XOR	−	FF	F	1	12820	402
13	AND	+	TT	T	0	11870	434
14	AND	+	TF	F	2	13160	540
15	AND	+	FF	F	3	13400	504
16	AND	−	TT	T	2	12810	579
17	AND	−	TF	F	2	12920	497
18	AND	−	FF	F	1	12540	438

Table 1: Mean reaction times and standard deviations (in miliseconds) from 30 runs of the simulations of each condition (ID), broken down by logical connective (LC), polarity (pol, + for affirmative, − for negative), truth values of propositions (prop) and of the sentence (sent), and the number of activated CoD nodes (nCoD).

Reaction time could potentially be used to assess the process account of the model. In the model, reaction time is estimated as the time elapsed between the moment when the P1 Intention node is activated and the moment when one of the sentence-level truth value nodes is activated. In the exemplary simulation, we showed how deciding that a relation cannot be grounded inherently takes longer than deciding that is grounded. We will now study this effect systematically.

Table 1 reports the mean reaction times produced by the model in 18 different conditions. Simulations of each condition were run 30 times. The model parameters were tuned to ensure that it made no errors. Reaction times vary across repetitions, however, due to the stochastic component of the neural dynamics Eq. 1. Reaction times are larger when affirmative propositions are false or negative propositions are true, but also when the compound sentence is false (we did not model negation at the sentence level). Both false affirmative and true negative propositions entail detecting that the relation cannot be grounded, as signaled by a CoD node. A false sentence is signaled by the failure of a truth node, stopping the default answer as signaled by the sentence CoD. We expect the number of activated CoD nodes to determine the reaction time of the model across conditions. In the table, the truth values of the proposition and of the sentence are listed together with the number of CoD activations involved.

To visualize this predicted dependence on the number of CoD activations, Figure 6 plots mean reaction times in increasing order while visualizing the number of activated CoD nodes through a color scheme. With one exception, the number of activated CoD nodes explains the change in reaction

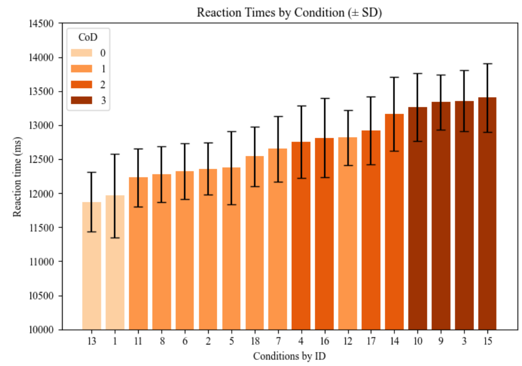

Figure 6: Mean reaction times in increasing order with error bar indicating standard deviation across repetitions. The conditions are labeled by the ID from Table 1. The color code reflects the number of activated CoD nodes in each condition

time across conditions. Given the complexity of the model, this regularity is informative and a useful basis for comparison to the experiment.

Discussion

The current work demonstrates a DFT-based model that determines the truth value of sentences that consist of two propositions conjoined by logical connectives OR, XOR, and AND. The propositions entail spatial relations and are perceptually grounded in a visual scene. A sequence of processing steps emerges autonomously from the underlying neural dynamics: Objects are sequentially attended and the match of their spatial configuration to the proposition's spatial relation is assessed. The truth value of each proposition is determined taking its polarity into account. Finally, the truth value of the compound sentence is determined by matching the truth values of the two propositions to the possibilities defined by the meaning of the logical connectives.

The propositions are represented by sustained activation in two copies of a set of neural nodes that represent the color feature of objects in two grammatical roles, target and reference, the spatial relation between these objects, and the polarity of the proposition. This is a simplified variant of how conceptual structure is represented in a DFT-based account of the perceptual grounding of nested phrases (Sabinasz & Schöner, 2023). In that work, the neural representation of object-in-role and relation concepts is enriched by index dimensions. As concepts are sequentially processed, their location along the index dimension is determined. The arguments of relations are bound through the shared index dimension. Arguments may be shared across different propositions, achieving through a slightly different mechanism a form of binding as proposed in neural network accounts for symbolic processing (Holyoak & Hummel, 2000). In the present model, only two propositions are invoked, corresponding to only two entries along the index dimension. The neural processes of perceptual grounding of the relations are a simplified over those in Sabinasz and

Schöner, 2023.

The treatment of polarity builds on Kati et al. (2024). The key idea is that relational propositions of negative polarity are evaluated by detecting the failure of perceptually grounding the affirmative version. In the model, this failure is signaled by the activation of the CoD node, the default response to grounding unless inhibited by the activation of the CoS node. Detection of failure is thus inherently slower than detection of successful grounding (see Chan and Hayward (2013) for analogous insights in visual search). This explains the negation-by-truth-value interaction effect (Kati et al., 2024) that has been observed psycho-physically (Kaup et al., 2005).

CoS and CoD nodes feed a neural representation of all possible combinations of polarity and truth value of the two propositions, which in turn drives the truth value determination of compound sentences. The compound sentences is evaluated as false by default, but that process may be interrupted by any of the truth-value combinations that vote for true. This is analogous to how a single spatial proposition is evaluated as false by default unless interrupted by the detection of grounding success, and leads to a similar signature in reaction time: The false outcome at the sentence level emerges on a slower time scale, increasing reaction time about as much as any activation of a CoD node does. The total number of CoD activation at both component and compound levels listed in Table1, explains the pattern of reaction times reported in Fig. 6.

Empirical assessment of such a difference in RT might face difficulty. Due to the complexity of the task and the large number of processing steps involved, RTs are expected to be large with concomitantly large variance. The predicted difference in RT may thus be a small portion of total variation and may not be detectable. There is also the question if the semantics of logical connectives in actual language usage are aligned with the assumption we made that all component propositions are evaluated. Lobina et al. (2023) experimentally estimated the time observers needed to verify compound sentences while looking at a visual scene. They found that verifying sentences conjoined by OR actually took longer than sentences conjoined AND, which might have to do with the scalar implicature relationship between these two connectives (Singh et al., 2016). This is consistent with our decision to exclude shortcut strategies in which observers would stop processing if one proposition of an *or*-sentence has been found to be true or, similarly, if one proposition of an *and*-sentence has been found to be false.

A main novelty of the model is the ability to explicitly represent multiple disjunctive possibilities given by the meaning of logical connectives. While Sabinasz and Schöner (2023) demonstrated how multiple nested phrases conjoined by *and* can be perceptually grounded, explicit representation of possibility was not needed as the *and* connective describes only one single possible world. That earlier work also did not preform truth value assessment. In contrast, the current work required the neural infrastructure to represent multiple possibilities implicated by disjunctives and to detect the match between such possibilities and the actual state of the world.

The current work is part of an effort to link accounts of perceptually grounded cognition to higher cognition by invoking the same principles of neural processing. The model builds on earlier processes models framed within DFT that include models of visual attention (Grieben et al., 2020), spatial relation grounding (Richter et al., 2021), and negation processing (Kati et al., 2024). The principles of DFT originate from embodied forms of cognition that are tightly coupled to sensory inputs, and potentially to motor outputs (Grieben et al., 2024). How such principles may also address more abstract forms of cognition is the classical challenge of grounded cognition (Barsalou, 2008).

Perhaps the most interesting extension of the current model to higher-order cognitive competencies would be to address reasoning based on compound sentences as premises in the absence of perceptual grounding. According to the Mental Model Theory (Johnson-Laird, 1983), humans reason by building iconic models that instantiate the relational structure given by the sentences. For example, when given multiple sentences describing spatial relations between objects, the spatial arrangements of objects may be represented as a spatial map (Ragni & Knauff, 2013). An earlier DFT model (Kounatidou et al., 2018) provided a neural process account for how such a spatial map may be built from the premises. The model then performs spatial inference by grounding a partial spatial phrase, the query, in the mental map. This grounding is achieved by the same kind of processes involved in perceptual grounding. The result of grounding the query is the activation of a location and feature representation of an object in the mental map that matches the query. This activates a corresponding conceptual description which provides the answer to the query.

The "principle of truth" of mental model theory may provide guidance for how to extend the earlier work on mental mapping and the present work on compound sentences: "Mental models represent only those possibilities in which an assertion is true, and in each of these models they represent only those clauses in the assertion that are true in the relevant possibility" (reviewed in Johnson-Laird and Ragni (2019), p.6). In the present model, logical connectives only boost nodes representing possibilities that lead the compound sentence to be true. However, within each possibility representation, the model explicitly represents the truth values of propositions including false ones. Also, mental maps built according to Kounatidou et al. (2018) would plausibly not include objects whose existence is negated in a truthful proposition. Reasoning with compound sentences requires the capacity to generate and operate on multiple mental models that represent different possibilities (Johnson-Laird, 2006). That has not yet been achieved in the prior work. If this were achieved, it could be combined with the representations of truth values developed here to perform the required inference.

References

Barsalou, L. W. (2008). Grounded cognition. *Annual review of psychology*, *59*(1), 617–645.

Chan, L. K., & Hayward, W. G. (2013). Visual search. *Wiley Interdisciplinary Reviews: Cognitive Science*, *4*(4), 415–429.

Grieben, R., Sehring, S., Tekülve, J., Spencer, J. P., & Schöner, G. (2024). ROBOVERINE: A human-inspired neural robotic process model of active visual search and scene grammar in naturalistic environments. *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 11470–11477. https://doi.org/10.1109/IROS58592.2024.10801621

Grieben, R., Tekülve, J., Zibner, S. K. U., Lins, J., Schneegans, S., & Schöner, G. (2020). Scene memory and spatial inhibition in visual search: A neural dynamic process model and new experimental evidence. *Attention, Perception, & Psychophysics*, *82*(2), 775–798. https://doi.org/10.3758/s13414-019-01898-y

Holyoak, K. J., & Hummel, J. E. (2000). The proper treatment of symbols in a connectionist architecture. In *Cognitive dynamics: Conceptual and representational change in humans and machines* (pp. 229–263). Lawrence Erlbaum Associates Publishers.

Johnson-Laird, P. N. (2006). Mental Models, Sentential Reasoning, and Illusory Inferences. In C. Held, M. Knauff, & G. Vosgerau (Eds.), *Advances in Psychology* (pp. 27–51, Vol. 138). North-Holland.

Johnson-Laird, P. (1983). *Mental models: Towards a cognitive science of language, inference, and consciousness*. Harvard University Press.

Johnson-Laird, P., & Ragni, M. (2019). Possibilities as the foundation of reasoning. *Cognition*, *193*, 103950.

Kati, L., Sabinasz, D., Schöner, G., & Kaup, B. (2024). Interaction of polarity and truth value-a neural dynamic architecture of negation processing. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *46*.

Kaup, B., Ludtke, J., & Zwaan, R. A. (2005). Effects of negation, truth value, and delay on picture recognition after reading affirmative and negative sentences. *Proceedings of the annual meeting of the Cognitive Science Society*, *27*(27).

Kounatidou, P., Richter, M., & Schoner, G. (2018). A neural dynamic architecture that autonomously builds mental models. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *40*.

Lobina, D. J., Demestre, J., García-Albea, J. E., & Guasch, M. (2023). Default meanings: Language's logical connectives between comprehension and reasoning. *Linguistics and Philosophy*, *46*(1), 135–168.

Lomp, O., Richter, M., Zibner, S. K., & Schöner, G. (2016). Developing dynamic field theory architectures for embodied cognitive systems with cedar. *Frontiers in neurorobotics*, *10*, 14.

Ragni, M., & Knauff, M. (2013). A theory and a computational model of spatial reasoning with preferred mental models. *Psychological Review*, *120*, 561–588.

Richter, M., Lins, J., & Schöner, G. (2021). A neural dynamic model of the perceptual grounding of spatial and movement relations. *Cognitive Science*, *45*(10), e13045.

Sabinasz, D., & Schöner, G. (2023). A neural dynamic model perceptually grounds nested noun phrases. *Topics in Cognitive Science*, *15*(2), 274–289.

Schöner, G., Spencer, J., Group, D. R., et al. (2016). *Dynamic thinking: A primer on dynamic field theory*. Oxford University Press.

Singh, R., Wexler, K., Astle-Rahim, A., Kamawar, D., & Fox, D. (2016). Children interpret disjunction as conjunction: Consequences for theories of implicature and child development. *Natural Language Semantics*, *24*(4), 305–352. https://doi.org/10.1007/s11050-016-9126-3