Meta-Analysis of the Functional Neuroimaging Literature with Probabilistic Logic Programming

Majd Abdallah^{1,2,†,*}, Valentin Iovene^{1,2,†}, Gaston Zanitti^{1,2}, and Demian Wassermann^{1,2,*}

¹Team Parietal, Inria, CEA, Université Paris-Saclay, Palaiseau, 91120, France

²NeuroSpin, CEA, Université Paris-Saclay, Gif-sur-Yvette, 91190, France

*Corresponding authors

†These authors contributed equally

February 25, 2022

Abstract

Inferring reliable brain-behavior associations requires synthesizing evidence from thousands of functional neuroimaging studies through meta-analysis. However, existing meta-analysis tools are limited to investigating simple neuroscience concepts and expressing a restricted range of questions. Here, we expand the scope of neuroimaging meta-analysis by designing NeuroLang: a domain-specific language to express and test hypotheses using probabilistic first-order logic programming. By leveraging formalisms found at the crossroads of artificial intelligence and knowledge representation, NeuroLang provides the expressivity to address a larger repertoire of hypotheses in a meta-analysis, while seamlessly modelling the uncertainty inherent to neuroimaging data. We demonstrate the language’s capabilities in conducting comprehensive neuroimaging meta-analysis through use-case examples that address questions of structure-function associations. Specifically, we infer the specific functional roles of three canonical brain networks, support the role of the visual word-form area in visuospatial attention, and investigate the heterogeneous organization of the fronto-parietal control network.

Keywords— Neuroimaging, Meta-Analysis, Probabilistic logic programming, Domain-specific language, Neuroinformatics, Knowledge Representation

1 Introduction

Meta-analysis confronts hypotheses of brain-behavior mappings with synthesized evidence from thousands of functional neuroimaging studies [1]. However, commonly used standard tools for meta-analysis, such as BrainMap [2] and Neurosynth [3] are limited in their *formal expressivity*. For example, it is challenging to express complex questions like ‘which topics are likely to be present in a study given activation in one set of regions and *there exists no* activation in another set of regions’. Answering this question requires that we explicitly run a meta-analysis on every single term of interest, while integrating heterogeneous and uncertain data, such as spatial masks and parcellations, to define the regions [4–7]. Importantly, however, applying the nontrivial criteria to select articles for inclusion and exclusion necessitates coding in a general-purpose programming language, which can be error-prone and hard to maintain. In this work, we harness advancements in symbolic artificial intelligence, specifically in *probabilistic logic programming*, to design *NeuroLang*: a domain-specific language (DSL) to formulate rich and expressive neuroscience hypotheses using formal logic-based criteria, while seamlessly combining heterogeneous data and reasoning about uncertainty. Our bet is that NeuroLang’s clear syntax and probabilistic logic semantics will enable accessible, rigorous, and highly reproducible large-scale meta-analysis.

Large-scale meta-analysis has become popular in cognitive neuroscience research as a result of the growth in non-invasive neuroimaging experiments, which provoked an upsurge in the number of yearly publications on structure-function associations [8]. As a powerful approach to synthesizing large quantities of results, meta-analysis has been performed to distinguish spurious from replicable findings, derive coactivation patterns [9], and define a priori regions of interest [10]. Moreover, elaborate coactivation-based connectivity, meta-analytic functional parcellations, and sophisticated network models are now being inferred from datasets covering tens of thousands of subjects. Recently, even more elaborate meta-analyses have been performed to derive new findings. For instance, Yeo et al. [6] have used meta-analysis to study the human association cortex, identifying frontal and parietal regions that are either specialized or flexible. Furthermore, meta-analysis has been used to support findings on the fractionation of networks into subsystems underlying disparate processing domains [7, 11]. Thus, large-scale meta-analysis is essential for the continued detection of latent properties of brain systems, far beyond what can be inferred from individual studies.

Over three decades ago, spatial normalization—the use of peak activation coordinates standardized in stereotaxic space—was introduced to human neuroimaging research [12]. This norm has been embraced, enlarged, and popularized by an outstanding series of methodological breakthroughs in the field. As a result, a large and methodologically cohesive functional neuroimaging literature has emerged over the past years. Capitalizing on this wealth of results, large-scale online databases have been created to compile peak activation coordinates and their related meta-data, spurring the development of coordinate-based meta-analytic approaches to exploit this rapidly growing corpus. This makes coordinate-based meta-analysis (CBMA) the most popular type of neuroimaging meta-analysis so far. The earliest approach to compiling peak coordinates, used by BrainMap [2], is to manually transcribe them from study tables. The peaks are then linked to the content of the studies by taxonomy experts. A more recent approach, used by Neurosynth [3], is to automate this laborious task via special software that automatically extracts peak coordinates from study tables. Neurosynth also automates the annotation

of studies using natural language processing techniques that incorporate term-frequency features, i.e. TF-IDF, estimated from the texts of studies. Albeit noisier, these fully automated approaches to meta-analysis scale better to the rapidly expanding neuroimaging literature.

Tools like BrainMap [2] and Neurosynth [3] have indeed simplified meta-analysis, becoming cornerstones of contemporary neuroscience research. However, only a narrow range of questions can be *natively* posed using these tools due to their limited *expressivity*. Specifically, these tools are based on propositional logic, which makes formulating queries beyond straightforward propositions, while applying nontrivial selection criteria of studies, a verbose and arduous task. More recently, NeuroQuery, a regularised predictive model trained on CBMA data, has been introduced to map *any arbitrary text fragment* (i.e. a set of keywords of interest, dubbed ‘query’) to a brain activation map. This approach shows that modelling semantic relations across studies can produce meaningful statistical maps for terms that are rarely mentioned in the literature. Nonetheless, NeuroQuery cannot express questions that infer a pattern of activation from studies associated with ‘working memory’, for instance, *but not associated* with other related functions. Such an analysis could be interesting to infer unique neural substrates among closely-related mental functions. Yet, NeuroQuery does not employ any *logic-based semantics*. That is, the words ‘but’ and ‘not’ are not in NeuroQuery’s pre-selected vocabulary, and have no influence on the resulting statistical maps. Thus, there seems to be a space for a logic-based framework that, by giving access to more elaborate meta-analytic queries, could help bridge the gap between statistical modeling and cognitive neuroscience. An all encompassing domain-specific language (DSL) for neuroscience research could encode any type of data, express complex queries using logic semantics, and reason about elements of uncertainty, all in a single unifying framework.

Parallel to the growth of functional neuroimaging over the past few decades, the artificial intelligence (AI) sub-fields of *probabilistic logic programming* and *probabilistic databases* have experienced a rapid evolution. This has led to the development of methods that efficiently represent knowledge with both *logical* and *probabilistic* semantics in a way that makes statistical model assumptions declarative, clear and less biased. Theoretical formalisms ground probabilistic logic languages with well-defined semantics and a mathematical understanding of the complexity of various classes of queries and different types of data [13]. Efforts to address the scalability of probabilistic logic systems led to recent algorithms and data structures that can remarkably speed up analyses, even on very large and diverse databases [14].

Can probabilistic logic programming formalize and broaden the range of questions that can be expressed in a meta-analysis? Can it simplify formulating complex hypotheses while combining heterogeneous and uncertain data? In this work, we present NeuroLang, a domain-specific language for conducting comprehensive neuroimaging meta-analyses using formal, succinct, and self-contained programs. At its core, NeuroLang uses predicate logic, as opposed to propositional logic, to ease the formulation of queries in a way that is closer to human discourse [15], and that can be run against structured neuroscientific data. Logic rules and queries are augmented with probabilistic semantics to account for uncertainty that emerges from missing information, analytical variability across studies, and measurement imperfections. In this article, we present concrete use-case applications of NeuroLang that address popular questions from the literature.

2 Results

The use-case examples of NeuroLang shed light on the utility of probabilistic logic semantics in representing neuroscience hypotheses that cannot be readily expressed with standard meta-analysis tools. In the first example, we use between-network segregation queries to infer the unique functional roles of three canonical functional networks: the dorsal attention network (DAN), default mode network (DMN), and frontoparietal cognitive control network (FPCN). In the second example, we explore potential associations between topics and activity within the visual word-form area (VWFA), when it either coactivates with regions of the dorsal attention network or those of the language network. In the third and fourth examples, we study the functional heterogeneity of the FPCN, uncovering differential activation profiles for a number of mental functions and varying connectivity patterns with other brain networks.

Representing Neuroscientific Knowledge under Uncertainty

Before exploring use-case examples of NeuroLang, we describe how heterogeneous neuroimaging data are represented in using fact and rule *tables*. A table is a set of tuples or rows, each representing a data instance and has a set of k elements representing columns. *Probabilities* can be ascribed to the rows of a table to quantify the level uncertainty in the data presented by each, in which case the table is said to be *probabilistic*.

Studies in a CBMA database report a set of *peak* activations that we store in a table named PeakReported. This table contains one row (x, y, z, s) for each peak that a study s has reported at location (x, y, z) in MNI space. Moreover, the uncertainty around the spatial location of peaks can be encoded in a rule table by assuming each peak’s 10mm neighboring voxels to be equivalently reported, similar to the multilevel kernel density analysis (MKDA) [16]. This rule table is called VoxelReported, and it includes a row (x, y, z, s) for each voxel at location (x, y, z) in the neighborhood ($< 10mm$) of a peak reported by study s . More details on how other spatial smoothing priors can be encoded in NeuroLang, such as the probabilistic prior used by the activation likelihood estimation (ALE) [17] algorithm, are provided in the Supplementary Materials.

Further, each study within a meta-analytic corpus is associated with cognitive processes or concepts addressed by its experiments. Fully automated meta-analytic tools like Neurosynth calculate statistical term-frequency features on study texts or abstracts, and threshold them to establish these links in a data-driven manner [3]. We store these associations within a TermAssociation table, containing one row (t, s) for each term t associated with study s . Moreover, we incorporate data-driven topic models, learned and openly shared by Neurosynth [18], within a TopicAssociation probabilistic table, containing one row $(t, s, \mathbf{P})$ for each uncertain association between a topic t and a study s . In probabilistic logic, we write $\text{TopicAssociation}(t, s) :: \mathbf{P}$ to state ‘study s has a probability $\mathbf{P}$ of being associated with topic t ’ [19]. This data representation process is illustrated in fig. 1.

Similarly to Neurosynth, we assume each study within the meta-analytic database to be an *independent equiprobable sample* of neuroscientific knowledge [3, 17]. This assumption is encoded by a SelectedStudy rule, depicted in the bottom left part of fig. 1, which gives studies an equal weight $1/N$ in any meta-analysis, where N is the total number of studies within the meta-analytic database. This makes it possible to estimate statistics on CBMA databases in the absence of statistical power indicators (e.g. sample size).

It is common for meta-analyses to integrate anatomical or functional brain parcellations [11] to enhance interpretability and reduce computational burdens. In the examples, we use the DiFuMo-256 atlas [20], which is part of a multiscale “soft” parcellation estimated from thousands of subjects across 27 studies that include both task and resting-state fMRI experiments. This data-driven functional atlas is argued to achieve comparable statistical performance as voxel-level analyses, while simultaneously reducing computational cost and enhancing interpretability. We represent the 256 functional regions from DiFuMo in a RegionVoxel table, containing a row (r, x, y, z) for each brain voxel at MNI location (x, y, z) belonging to a DiFuMo-256 region r . An excerpt of the RegionVoxel table is depicted on the right part of fig. 1. We also incorporate a table NetworkRegion that contains a row (n, r) for each region r that significantly overlaps with some network n from the 7 or 17-network parcellations [21]. The network membership of regions is provided as part of the DiFuMo meta-data file. The experiments that follow will use this unified framework of knowledge representation to formulate probabilistic logic programs that drive meta-analytical findings.

Figure 1: Representation of meta-analytic and functional parcellation knowledge using database tables in NeuroLang.

Between-Network Segregation: Reverse Inference of Brain Network Function

In this example, we perform a *segregation-based meta-analysis* to infer the likelihood of a topic to be present in a study given activation in a brain network, with an additional constraint that there *exists no activation in other networks*. The goal of this example is to show that a segregation query can identify which network’s activation pattern is preferentially more predictive of the presence of topic terms related to certain mental functions.

We use the Neurosynth CBMA database [3], consisting of 14,371 studies, and its associated v5-topics-100 topic model [18]. The networks included in this example are the DMN, FPCN and DAN defined using the coarse 7-Network atlas [21]. These networks exhibit coupling dynamics in support of an array of internally and externally-

directed mental functions [22]. However, each one of them is believed to subserve a unique set of cognitive processes [7, 22, 23]. The FPCN contributes to a wide variety of tasks by engaging top-down control processes, the DAN is concerned with orienting attention towards salient cues, and the DMN is involved in abstract self-referential, social and affective functions. Using a segregation query, we can quantitatively identify the specific functional roles of these networks from the literature.

First, we have to represent useful heterogeneous data in NeuroLang. For instance, we assume a DiFuMo-256 component r to be reported by a study s whenever a peak activation is reported by the study within that region. In NeuroLang, this is expressed by the following logic rule

```
RegionReported( $r$ ,  $s$ ) :-  
    PeakReported( $x$ ,  $y$ ,  $z$ ,  $s$ ) & RegionVoxel( $r$ ,  $x$ ,  $y$ ,  $z$ )
```

which translates, in plain English, to ‘region r is reported by study s if s reports a peak at location (x, y, z) that falls within region r ’. Furthermore, we model the reporting of networks by studies in a *probabilistic* table. The probabilities are based on the total volume of the reported regions that belong to a network. This table accounts for the uncertainty in the location of reported peak activation coordinates as well as the number of potentially reported regions. More precisely, we consider that each study has a probability of reporting a network, proportional to the number of reported regions belonging to the network.

This is implemented by the following rules in NeuroLang

```
RegionVolume( $r$ , count( $x$ ,  $y$ ,  $z$ )) :-  
    RegionVoxel( $r$ ,  $x$ ,  $y$ ,  $z$ )  
  
NetworkVolume( $n$ , sum( $v$ )) :-  
    RegionVolume( $r$ ,  $v$ ) & NetworkRegion( $n$ ,  $r$ )  
  
ReportedVolume( $n$ ,  $s$ , sum( $v$ )) :-  
    RegionVolume( $r$ ,  $v$ ) & NetworkRegion( $n$ ,  $r$ )  
    & RegionReported( $r$ ,  $s$ )  
  
NetworkReported( $n$ ,  $s$ ) :  $v/V$  :-  
    ReportedVolume( $n$ ,  $s$ ,  $v$ ) & NetworkVolume( $n$ ,  $V$ )
```

In plain English, ‘a network n is considered to be reported by study s with probability v/V , where v is the total volume of regions within network n that are reported active by study s , and V is the total volume of all regions in the network’. This program makes use of NeuroLang’s built-in `count` and `sum` aggregation functions.

Next, we define a rule that infers the probability that studies are associated with a topic *given* activation in only one of the three networks. That is, we query the probability that a study s is associated with topic t given that some network n is reported by s and *there exists* no other network reported by study s . In NeuroLang, this corresponds to the following rule

```
ans(t, n, PROB) :-
  TopicAssociation(t, s) & SelectedStudy(s)
  // NetworkReported(n, s) & ~exists(n2;
    Network(n2) & n2 != n & NetworkReported(n2, s)
  )
```

where the `//` operator is read as *given*, representing *probabilistic conditioning*. This rule contains a negated existential expression, `~exists(...)`, that prevents two or more networks from being reported by a study at the same time.

We report the resulting functional profiles in fig. 2. We observe that topics related to sensory processing of direct environmental demands such as eye movements, visual attention, and spatial orientation are more likely to appear in studies reporting activations in the DAN only. Also, we observe that topics related to cognitive control such as task switching, task demands, response inhibition, and performance monitoring are more likely to be mentioned in studies reporting activations in the FPCN. Finally, topics related to higher-order abstract cognitive and memory-related processes are mostly associated with studies reporting DMN activations only.

Figure 2: Functional profiles obtained with network-based segregation queries that identify the most probable topic associations in studies reporting activations within one network but not reporting activations within any of the other networks. A 95% confidence interval is depicted, across 1000 random 50% subsamples of the Neurosynth database [3].

Meta-Analysing the Role of the Visual Word-Form Area in Attention Circuitry

The visual word-form area (VWFA) has attracted controversy over the years with recent findings suggesting it to take part in the attention circuitry not only in the language network [24]. Can this relationship be inferred solely from a meta-analysis of past studies that have reported activations in the left ventral occipito-temporal cortex without necessarily identifying it as the VWFA?

To answer this question, we formulate queries that infer the most probable topic associations among studies that report activations close to the VWFA region, while simultaneously reporting activations within regions of the attention network, but *not* reporting activations within regions of the language network.

To define regions of the VWFA; dorsal attention and language network, we use the seed locations defined in [24], and store them in a `RegionSeedVoxel` table. This table contains a row (x, y, z, r) for each region r 's seed location (x, y, z) . A database table `NetworkRegion` contains rows (n, r) for each region r belonging to network n . A brain region is considered to be reported by a study if it reports a peak activation within 10mm of the region seed location. The choice of a 10mm radius was used to facilitate comparisons with the range of smoothing kernels that are typically used within meta-analyses.

In this example, a network is considered to be reported by a study if it reports one of the network's regions. In NeuroLang, these rules are as follows:

```
RegionReported(r, s) :-
    PeakReported(x, y, z, s) & RegionSeedVoxel(x1, y1, z1, r)
    & d = EUCLIDEAN(x, y, z, x1, y1, z1) & d < 10
NetworkReported(n, s) :- RegionReported(r, s) & NetworkRegion(n, r)
```

where `EUCLIDEAN` is a built-in function that calculates the Euclidean distance between two locations in MNI space.

Finally, to test our hypothesis, we use the following *probability encoding rule*

```
ans(t, n, PROB) :- TermAssociation(t, s) //
    (RegionReported(vwfa, s)
    & NetworkReported(n, s)
    & SelectedStudy(s)
    & ~exists(n2; Network(n2) & n2 != n & Study(s)
    & NetworkReported(n2, s))
    )
```

which calculates the probability of finding an association with topic t among studies that report the activation of both the VWFA and network n , but do not report the activation of any other network n_2 , where $n_2 \neq n$. Because only two networks, `language` and `attention`, are present in the `Network` table, this rule simultaneously calculates the probabilities for each pair of networks, including one while segregating the other.

Results are shown in table 1. Topic 32 was found to be significantly associated with studies that report activations within the VWFA and the attention network but that do not report activations within the 'language' network. This topic loads on terms related to *object recognition*—a task for which attention circuitry is essential [25]. This result suggests that the VWFA may play a role in attention, as studies that report its activations

Topic	Likelihood Ratio Test
32.object_objects_visual	$\chi^2(1, N = 455) = 15.96, p_{\text{FDR}} = 0.0065$
21.reading_words_word	$\chi^2(1, N = 455) = 14.4, p_{\text{FDR}} = 0.0074$

Table 1: Topics associated with studies reporting the VWFA and the fronto-parietal attention network, but not reporting the ‘language’ network. We only depict associations surviving a likelihood ratio test and a false discovery rate (FDR) correction for multiple comparison ($p_{\text{FDR}} < 0.05$). More details on the likelihood ratio test are provided in the Supplementary Materials. Topics are those from Neurosynth’s `v5-topics-100` topic model. Out of Neurosynth’s 14,371 studies, 455 report activations within the fronto-parietal attention network without reporting activations within the ‘language’ network.

are significantly associated with object recognition, and supports the running hypothesis that the VWFA plays a role in processing multiple categories of visual stimuli [24]. We also observe a significant association with topic 21, which loads on terms related to the task of *reading words*—the putative role of the VWFA.

The *opposite* segregation query selects studies reporting the VWFA and the ‘language’ network but not reporting the attention network ($N = 318$). This analysis did not yield any significant topic association after correction for multiple comparison. However, a similar topic association analysis, but without segregating studies that report activation in the attention network, does yield a significant association with topic 21, linked to the ‘reading words’ ($\chi^2(1, N = 852) = 56.86, p_{\text{FDR}} = 0.000081$). This result might have more than one explanation, but a plausible explanation could be the relative decrease in statistical power (i.e. smaller number of studies) in the segregation query compared to the non-segregation query.

Inferring Differential Activation Patterns within the FPCN using Topic Segregation Queries

In this example, we perform forward inference using *topic-based segregation queries* to derive activation patterns within the frontoparietal cognitive control network (FPCN). As a major part of the *multiple demand* system [26], the FPCN is associated with a large set of tasks, themselves belonging to disparate and overlapping cognitive processes such as working memory, memory retrieval, task switching, and semantic processing, to name a few. Moreover, there is evidence for a heterogeneous internal organization in the FPCN, whereby a different combination of regions may be involved in a different domain of control processing [27]. Thus, the goal of this example is to infer activation patterns within the FPCN predicted by the presence of topic terms related to one process and the simultaneous absence of topic terms related to other processes. In this sense, segregation queries can enhance the relative specificity of meta-analytic forward inferences by minimizing the amount of overlap amongst related topics.

From the set of 200 Neurosynth topics (version-5), we select five exemplar topics representing a subset of the cognitive processes often attributed to the FPCN, along with the loading values of studies on each topic. These topics are *working memory*, *decision making*, *task set switching*, *semantic control*, and *memory retrieval* [22, 26, 28, 29]. Then, we formulate the following NeuroLang program which performs topic segregation queries, yielding an activation map for each topic separately. This program is written as follows:

Figure 3: Cortical maps showing the difference in posterior probabilities of FPCN regions to be active given topic segregation and when given no topic segregation queries. We mask out brain voxels that are not part of the FPCN. The difference between posterior probabilities is defined as $\Delta = \mathbf{P}[\text{VoxelReported}(x, y, z) | \text{SingleTopicAssociation}(t)] - \mathbf{P}[\text{VoxelReported}(x, y, z) | \text{TopicAssociation}(t)]$. Probability differences are thresholded at 0.05.

```
SingleTopicAssociation(topic, s) :-
  TopicAssociation(topic, s) & ~exists(t;
    Topic(t) & t != topic & TopicAssociation(t, s)
  )
ans(topic, x, y, z, PROB) :-
  VoxelReported(x, y, z, s) & SelectedStudy(s)
  // SingleTopicAssociation(topic, s) & SelectedStudy(s)
```

We report the resulting topic-based activations within the FPCN in fig. 3. The results of this segregation query show that the FPCN exhibits a varied activation profile across topics, corroborating previous findings of flexible adaptation of activity within this network as task demands change. Specifically, working memory and task set switching tend to activate, to some extent, spatially interleaved, frontal and parietal regions of the FPCN network. Semantic processing, on the other hand, dominantly activates a left-lateralized ventral frontal regions. Finally, decision making and memory retrieval are associated with activation in the cingulo-medial portion of the FPCN, the pre-supplementary motor/dorsal anterior cingulate cortex (decision making) and a precuneus/posterior cingulate cortex network (memory retrieval).

Inferring Varying Meta-Analytic Connectivity Profiles of FPCN Sub-networks

Recent findings suggest that the frontoparietal cognitive control network (FPCN) can be decomposed into sub-systems associated with disparate and overlapping mental processes. Dixon et al. [7] studied two broad subsystems of the FPCN that also appear as separate networks in the influential 17-network model from Yeo et al. [21]. Using the same nomenclature, we label these two subsystems FPCN-A and FPCN-B. Dixon et al. observed preferential connectivity between FPCN-A and the default mode network (DMN), and between FPCN-B and the dorsal attention network (DAN). We reproduce these results by conducting a similar, but more compact, meta-analysis with NeuroLang.

For this analysis we use the NeuroQuery [30] database instead of Neurosynth. We formulate conditional probabilistic queries that include studies reporting activations in each of the two FPCN sub-networks. By contrasting their posterior probability maps, we identify a distinct meta-analytic connectivity pattern associated with each sub-network. Using the same probabilistic definition of network reported by studies as in the first example, we formulate a rule that calculates the coactivation pattern of each FPCN sub-network. In NeuroLang, we use the following rule to calculate the conditional probability of a region being reported given that a network is also reported

```
ans(r, n, PROB) :- RegionReported(r, s) //
  (NetworkReported(n, s) & SelectedStudy(s))
```

whose resulting `ans` table contains rows (r, n, p) , where p is the probability of region r being reported active given that network n , where n is either FPCN-A or FPCN-B.

A likelihood-ratio test and a FDR correction ($\alpha = 0.05$) for multiple comparison are used to identify statistically significant coactivating regions. To ensure that the results are not driven by one choice of studies, we estimate the conditional probabilities on 1000 random sub-samples of the NeuroQuery database (each sub-sample is 50% of the entire database). This allows us to compute confidence intervals for our probability estimates,

Figure 4: Comparison of the probabilities that DiFuMo-256 components coactivate with the two FPCN subnetworks. Regions are colored based on their belonging to the networks proposed by Yeo et al. [21]. Only regions exhibiting a statistically significant ($p_{FDR} < 0.05$) coactivations with either subnetwork are included in the figure, based on the likelihood-ratio test and a correction for multiple comparison. $\mathbf{P}[\text{RegionReported}(r) | \text{NetworkReported}(\text{FPCN-A})]$ denotes the conditional probability of region r being reported by studies reporting FPCN-A in the database. Probabilities are calculated on 1000 random 50% subsamples of the NeuroQuery CBMA database.

which reflect the amount of variance across sub-samples.

In fig. 4, we show scatter plots of the probabilities that each DiFuMo-256 brain region is active given activation of the FPCN-A or FPCN-B sub-networks defined by Yeo 17-network parcellation. We only show the results of regions that exhibit a significant coactivation with at least one subnetwork, based on a likelihood-ratio test. In the left panel, regions are color coded by their network membership according to the coarser Yeo 7-network parcellation to facilitate interpretation.

Figure 5: DiFuMo-256 components that are more likely to coactivate with one FPCN sub-network than the other. In blue, we depict regions exhibiting a greater probability of coactivation with FPCN-A. In red, we depict regions exhibiting a greater probability of coactivation with FPCN-B. The absolute difference between region coactivation probabilities is defined as $\Delta = \mathbf{P}[\text{RegionReported}(r)|\text{NetworkReported}(\text{FPCN-A})] - \mathbf{P}[\text{RegionReported}(r)|\text{NetworkReported}(\text{FPCN-B})]$. A likelihood-ratio test and a FDR correction ($\alpha = 0.05$) for multiple comparison are used to identify regions that exhibit significant coactivation with either networks before estimating Δ .

In general, regions belonging to the somatomotor, visual, salience networks do not preferential coactivate with either the FPCN-A or FPCN-B. In contrast, regions of the coarse FPCN show a dichotomy in their coactivations with either FPCN-A and FPCN-B. That is, meta-analysis supports the hypothesis that FPCN can be functionally divided into two sub-systems [7]. Importantly, we find a clearer dichotomy in the coactivation profiles of the DMN and the DAN with the FPCN sub-networks. On one hand, 31 out of 32 DMN regions coactivate more with FPCN-A, while only one DMN region (a sub-region

in the middle frontal gyrus) seem to exhibit a preferential coactivation with FPCN-B. In fig. 5, we illustrate a meta-analytic coactivation contrast map between FPCN-A and FPCN-B, showing that the former coactivates to a greater extent with the core regions of the DMN, than does the latter. On the other hand, without indicating any preference, we observe that 21 out of 30 DAN regions exhibit statistically significant coactivations with FPCN-A, while 19 show significant coactivations with FPCN-B. However, only 11 regions have a higher probability of activating given an FPCN-B activation than FPCN-A, while the others have comparable probabilities of coactivating with either sub-networks. This is in line with the findings from Dixon et al. [7], showing less distinction in the DAN with respect to coactivation with the FPCN sub-networks. Nonetheless, FPCN-B does indeed coactivate to a greater extent with the core regions of the DAN, the superior parietal lobule and frontal eye fields, than FPCN-A as seen from the coactivation contrast map in fig. 5.

3 Discussion

We present a domain-specific language (DSL), coined NeuroLang, which broadens the range of meta-analytic hypotheses that can be expressed and tested against an ever-increasing functional neuroimaging literature. Through probabilistic logic semantics, users can formally represent their hypotheses, query heterogeneous data, and reason about uncertainty in a unified language. Ultimately, NeuroLang is envisioned to lead a new generation of computational tools for neuroimaging data analysis, including meta-analysis, to reduce miscommunication in the community and promote formal and reproducible research. To support this language-oriented approach, we provide concrete meta-analysis examples fully performed with NeuroLang.

An important meta-analysis application, beyond finding consistent activation patterns, is inferring reliable and *specific* structure-function associations. Traditionally, researchers carefully select studies that "ask a similar question" or rely on databases of expertly curated annotations of studies, such as BrainMap [2]. However, non-automated meta-analysis is not scalable, time-consuming, and can suffer from low statistical power. As a response, automated tools have been developed, like Neurosynth [3], NeuroQuery [30], and more recently NiMare, to enable scalable, richer, and unbiased meta-analysis. However, as mentioned before, these tools cannot formally express nontrivial inclusion/exclusion criteria of studies to infer specific structure-function associations. In contrast, NeuroLang provides a unified formal framework to succinctly express versatile queries of functional specificity in the brain via *first-order logic semantics* rather than propositional logic statements [19].

A recurring use-case of expressive querying throughout the examples is that of *segregation queries*, which we formulate using the first-order logical negation operator ($\neg$) and existential quantifier ($\exists$). Using a segregation query enables a seamless split-up of studies in a meta-analysis, while contrasting any number of topics and brain regions of interest both in a forward and reverse inference paradigms. In this sense, segregation queries can enhance the specificity of inferred structure-function associations in brain regions that are putatively recruited to varying degrees by multiple tasks and brain networks, such as the VWFA [24] or the anterior insula [31] for instance.

The three networks example serves as a proof of concept for using segregation queries to derive specific structure-function associations for brain structures, such as networks.

In this example, we infer the topics that are preferentially linked to each of the FPCN, DMN, and DAN. These three networks are known to exhibit competitive and cooperative coupling dynamics across a wide array of tasks [32]. Thus, observed activations may emerge from the dynamics of these networks depending on task demands [33]. This might yield to a blurry characterization of networks in terms of their specific roles and spatial arrangement, sometimes leading to nomenclature ambiguities across studies and research groups [34]. One way to solve this is to infer relatively specific functions of a network by "isolating" its activity pattern from the other networks. Achieving this isolation through segregation queries, we find preferential topic associations for the FPCN, DMN, and DAN that align with the general understanding of their roles [26, 35, 36]. Such a segregation-based meta-analysis can be performed in future studies, for instance, to create a foundation for a fine-scale taxonomy of brain networks depending on their inferred roles in addition to their connectivity patterns.

In the VWFA example, we use segregation queries to address a controversial hypotheses [24, 37] about the role of this region in general visual processing beyond reading words. The main result is a statistically significant association for the VWFA with a topic loading on terms of object recognition given coactivation with attention but not language regions. However, the topic loads on terms such as 'fusiform' and 'occipitotemporal'—zones in the close vicinity of the putative VWFA known to represent features and attributes of objects [38], which might bias the results due to uncertainty in precisely locating the VWFA and the ambiguity in nomenclature across studies. But given that the dorsal attention network is active and its strong functional connectivity to the left occipito-temporal zone [37], a link between the VWFA and "object recognition" is plausible. Concurrently, we observe no statistically significant topic associations for the VWFA when it coactivates with language but not attention regions. Yet, if we do not exclude studies reporting the coactivity among language and attention networks, a significant association with a topic related to "reading words" is observed. This finding suggests that the VWFA might not be specific to language per se, but rather have a broader role at the interface of language and visuospatial attention. Findings from [24] suggest that the VWFA acts as a gateway between attention and language networks, such that the former amplifies the representations of written words so they may be conveyed to the latter. Our findings suggest that the VWFA is a more general visual processor that may be recruited in other visual tasks in addition to reading. This segregation-based meta-analysis can be extended to understand the dynamic roles of flexible regions in the brain, such as connector hubs.

In the third and fourth examples, we derive varying coactivation patterns within the FPCN revealing its heterogeneous organization. The FPCN comprises regions that coactivate across diverse conditions. Given the lack of formal definitions and fine lines between different executive functions, which are often conjointly studied, it can be difficult to determine domain-specific FPCN regions. As a potential solution to this problem, a segregation query can simultaneously select studies highly loading on topic terms related to a single function, discard studies loading on topic terms of other functions, and contrast them. The results reveal a relatively unique coactivation pattern *consistently* associated with each topic, consistent with findings of dynamical activity in canonical brain networks as a function of varying demands [39]. We say "relatively unique" because we only study a small set of topics for the sake of demonstration, while in fact there are putatively more functions attributed to the FPCN. This meta-analysis can include a cognitive ontology [40] to systematically define *all* pertinent concepts and similarly contrast them [41]. Moreover, this type of meta-analysis can be effective in system-level

causal modeling approaches (e.g., dynamic causal modeling) [42], which require strong a priori hypotheses about the regions involved in particular contexts. Finally, we reproduce the results of Dixon et al. [7], equivalently revealing two subsystems, FPCN-A and FPCN-B that exhibit distinct activations profiles and specific associations with the DMN and DAN, respectively. Although Dixon et al. [7] have successfully performed this analysis using Neurosynth’s command line tools, we have been able to reproduce their findings with significantly more compact, declarative and formal queries. In this sense, after representing the data in NeuroLang, a user only has to worry about what question to ask rather than about explicitly declaring every step needed to answer it.

Performing meta-analyses often requires integrating heterogeneous data. For example, Andrews et al. [11] use a brain parcellation to characterize components of the DMN whose respective functions are decoded through reverse inference reasoning. Here, to study the functional profiles of brain networks and the coactivation patterns of FPCN subsystems, we integrate the DiFuMo-256 functional atlas. Components of this functional atlas overlap with anatomical landmarks whose names are used to label the components by experts to enhance interpretability [20]. The DiFuMo components are also grouped into 7 and 17 canonical networks whose labels have been integrated in our examples. This approach facilitated the formulation of hypotheses and the interpretation of results. We believe that future studies conducted with NeuroLang could benefit from its capacity to represent any anatomical and functional atlas as well as tabular meta-data and formal ontologies [41]. Moreover, it is imperative for a *complete* meta-analytic tool to be flexible enough to represent any type of parcellation, meta-analytic database or, more generally, neuroscientific knowledge. Within our experiments, we have been able to represent both the Neurosynth database and its associated openly-shared topic models in NeuroLang. But, in other experiments, we use the NeuroQuery database because of its lower error-rate in the extraction of peak activation coordinates [30]. Together, these examples demonstrate that NeuroLang is agnostic to the database used for conducting meta-analyses, and could incorporate future sources of neuroscientific knowledge with various topologies.

Of course, due to the analytical variability across neuroimaging studies [43] and imperfections in data acquisition, knowledge representation should account for elements of uncertainty in the data. Probabilistic programs and databases constitute *general frameworks* for representing *structured* but *uncertain* knowledge. As these two paradigms reside at the heart of NeuroLang, uncertain data can be combined within its probabilistic programs. In our experiments, we model the reporting of functional modes and networks probabilistically based on their volumetric proportion that is reported by studies, although other indicators of uncertainty can be used. Another possible measure of uncertainty can be sample size and relative location of peaks with respect to the regions of a network. This probabilistic definition is thus *arbitrary*, but NeuroLang is expressive enough to represent any assumption just as well. For instance, to obtain network-specific functional profiles, we combine meta-analytic data and functional atlases with data-driven topic models. The topics are associated to studies probabilistically by the loadings of the fitted latent Dirichlet allocation model that are based on the frequency of co-occurrence of terms in abstracts of studies [18].

It is worth noting that, so far, not all types of queries can be *efficiently* solved by NeuroLang’s engine. For example, *negated* disjunctive queries, such as $\neg(\text{FirstSegregation}(s) \vee \text{SecondSegregation}(s))$, where FirstSegregation and SecondSegregation correspond to two segregation rules — similar to those presented in our examples —, can be computationally intractable in a voxel-level analysis, where $> 150,000$ voxels need to be modeled.

This is because extensions to logic programming, such as negation [44], are not directly transposable to lifted query processing of unions of conjunctive queries on *probabilistic* databases [45]. In other words, more expressivity in a language leads to higher *data complexity*—the complexity of solving a query with respect to the size of the data [46]. In the case of whole-brain voxel-level statistical modeling, data complexity is significant, and solving a segregation query can become impractical due to the large number of voxels that are modeled. Efficient algorithms exist for solving a limited *subset* of queries on probabilistic databases [14]. NeuroLang privileges *lifted query processing* to solve queries on probabilistic databases (see the Methods section for details). Lifted query processing is an algorithm with a set of rules that translate a query to an algebraic expression to solve it in polynomial time. This allows NeuroLang to efficiently scale to large databases. However, to calculate the solution of a potentially non-liftable query, NeuroLang falls back to knowledge compilation strategies [47]. Yet, when modeling hundreds of thousands of brain voxels, we find resolution times of knowledge compilation to be impractical as well. Obtaining a solution for meta-analytic queries at the voxel-level currently take several minutes to be solved with NeuroLang, while such queries can be solved in a few seconds by Neurosynth’s engine, which uses a custom python-based implementation. Nevertheless, major improvements to NeuroLang’s engine are currently underway.

In designing NeuroLang, we provide a high-level programming interface for harnessing meta-analysis databases in cognitive neuroscience research. We believe that this approach has *three* main advantages : Accessibility, Readability, and Sound semantics. Implementing a program to test a complex hypothesis against meta-analytic data can be time-consuming and error-prone, especially for those not proficient in general-purpose programming languages. A domain-specific syntax as NeuroLang’s eases the process of formulating hypotheses combining heterogeneous data. This has the potential of speeding up the meta-analysis process as well as making it highly reproducible. Moreover, being specific about the research question and assumptions in meta-analysis is an important practice [48]. NeuroLang’s logical syntax makes model assumptions and inclusion/exclusion criteria readable and understandable directly from the code of the program. Finally, NeuroLang is grounded in formal mathematical logic, providing theoretical guarantees that both limit modeling errors and provide trust in the language. We believe that these advantages make NeuroLang a tool of choice for conducting and reproducing functional neuroimaging meta-analyses. One NeuroLang program developed for a study in 2022 could be re-used by another team of researchers in 2027, five years later, either as a modeling inspiration for an entirely different study or to confront the original study’s findings with results published from 2022 to 2027.

4 Methods

We apply a language-oriented programming approach to the problem of expressing and testing meta-analytic neuroscientific hypotheses. That is, instead of using a general-purpose programming language to solve the problem, we design a domain-specific language (DSL) to represent the problem, and solve it in that language. NeuroLang uses *logic*, *declarative*, and *probabilistic* programming paradigms.

We start by introducing computer science concepts required to understand the semantics of the language. Then, we explore how these formalisms can be applied to the specific case of expressing cognitive neuroscience hypotheses. Finally, we detail the tech-

nicalities of solving queries on NeuroLang programs when working at the whole-brain neuroimaging scale.

Probabilistic Logic Programs and Databases

To *represent* heterogeneous neuroscientific knowledge, NeuroLang leans on Datalog [49], a fully declarative logic programming language designed to efficiently solve queries on large deductive databases. A Datalog rule takes the form $\forall \mathbf{x}, (\psi(\mathbf{x}) \leftarrow \exists \mathbf{y}, \varphi(\mathbf{x}, \mathbf{y}))$, where $\mathbf{x}$ is a set of universally quantified variables, $\mathbf{y}$ is a (possibly empty) set of existentially quantified variables, ψ is a relational symbol, and $\varphi(\mathbf{x}, \mathbf{y})$ is a conjunctive logic formula over $\mathbf{x} \cup \mathbf{y}$. For readability, the implicit $\forall$ and $\exists$ quantifiers are often omitted, and the rule is written $\psi(\mathbf{x}) \leftarrow \varphi(\mathbf{x}, \mathbf{y})$. A Datalog program is a set of such rules. The input of the program is a set of *extensional* facts and its output is a set of *intensional* facts that have been obtained through deductive inference, based on the program's rules and input facts. This process is summarised in fig. 6. In NeuroLang, we write a Datalog rule $\forall(x, z), (P(x, z) \leftarrow \exists y, Q(x, y, z) \wedge R(z))$ as $\mathbf{P}(\mathbf{x}, \mathbf{z}) \text{ :- } \mathbf{Q}(\mathbf{x}, \mathbf{y}, \mathbf{z}) \ \& \ \mathbf{R}(\mathbf{z})$. Theoretical formalisms and efficient query resolution algorithms have been developed by the logic programming community over the past decades. By restricting the syntax of its rules, Datalog queries have been proven to be solvable in polynomial time w.r.t. the size of the database [49]. Strong guarantees on query resolution complexity are primordial to handle high dimensional whole-brain neuroimaging data.

Figure 6: Deductive inference in Datalog. Using knowledge and implication rules to infer new knowledge.

Figure 7: A probabilistic logic program defines a probability distribution over its possible outcomes.

Testing cognitive neuroscience hypotheses requires aggregating data from many subjects or studies using statistical models. Logic programming languages were not designed

for statistics or probabilistic modeling, as their programs live in a world where everything is either true or false and only one outcome is possible. To address this limitation, logic programming languages were extended with probabilistic semantics to incorporate uncertainty and allow for probabilistic inference, such as ProbLog2 [50] or CP-Logic [51]. Following CP-Logic’s semantics, a probabilistic rule $\psi(\mathbf{x}) : \alpha \leftarrow \varphi(\mathbf{x})$ causes the event $\psi(\mathbf{x})$ to be true with probability α , whenever $\varphi(\mathbf{x})$ is true. Probabilistic logic programs define a distribution over possible outputs of the programs’ execution, a.k.a. *possible worlds* [52]. For a more detailed introduction to probabilistic logic programming, we refer the reader to De Raedt et al.’s review [53]. In NeuroLang, we write a CP-Logic rule $P(x, y) : f(x, y) \leftarrow Q(x, y) \wedge R(y)$ as $\mathbf{P}(\mathbf{x}, \mathbf{y}) : \mathbf{f}(\mathbf{x}, \mathbf{y}) :- \mathbf{Q}(\mathbf{x}, \mathbf{y}) \ \& \ \mathbf{R}(\mathbf{y})$.

Inferring the probability of a query boils down to summing the probabilities of all possible worlds where this query is verified. This process is called *weighted model counting* [54]. The number of possible worlds is often very large, and naively counting possible worlds becomes intractable. For example, were we to model the activation of each brain voxel as independent Bernoulli random variables, the number of possible worlds would be 2^K , where K is the number of voxels in the brain, typically hundreds of thousands. Solving weighed model counting problems on real-world data requires efficient resolution algorithms. Knowledge compilation finds compact representations of large probabilistic programs, and can be used to solve queries drastically faster [14, 55].

Figure 8: Solving a union of conjunctive queries on a tuple-independent database.

In parallel, the field of probabilistic databases extended traditional relational databases with the possibility of encoding uncertain knowledge using probabilities [56]. A probability can be attached to any tuple in the database, as illustrated in fig. 8. Similarly to probabilistic programs that define a distribution over possible outputs, a probabilistic database defines a probability distribution over a set of possible databases (or worlds), where tuples are chosen to be true or false based on their probability. The probability attached to a tuple then corresponds to the marginal probability of that tuple being found in any database randomly chosen from the distribution over possible databases. Probabilistic tuples within the database are often assumed to be independent random events, in which case the database is called a tuple-independent database (TID). A tuple with probability 1 is true in all possible databases, and a tuple with probability 0.5 is true in half of the possible databases.

In some cases, solving a query on a probabilistic database can be done much more efficiently than through knowledge compilation approaches. One recent theoretical result is the *dichotomy theorem*, which classifies queries on TIDs based on their complexity

Figure 9: An extensional query plan solves $\mathbf{P}[Q]$ for a given query Q by using algebraic operations that are extended with probabilistic calculations [14].

[13]. At the heart of this theorem, is a resolution strategy named *lifted query processing*. It applies, based solely on syntactic analysis of a queries, a set of rules that derive an algebraic expression that computes the probability of the query [14]. We illustrate this process in fig. 9. Lifiable queries have a polynomial data complexity. If the rules fail to apply, the query is said to be *non-lifiable*, and has been proven to have a #P-hard complexity. To solve non-lifiable queries, we use knowledge compilation strategies, which can be intractable in the voxel-level modeling setting.

By incorporating probabilistic semantics and fast query resolution algorithms from both probabilistic logic programming and probabilistic databases, NeuroLang is a full-fledged *probabilistic* programming language [19]. This approach makes it possible to express a wide variety of programs and queries, some of which can be efficiently solved using lifted query processing on probabilistic databases, even at the voxel-level.

Syntactic Specificities of NeuroLang

We describe the syntactic extensions of typical logic and probabilistic logic programming languages that we made to provide features necessary to express end-to-end meta-analyses in NeuroLang.

A NeuroLang probability encoding rule (PER) captures the result of a probabilistic inference into a deterministic table. They are a syntactic convenience, or *sugar syntax*, that makes it possible to solve probabilistic queries within the program and process their solution with deterministic rules. Internally, we *stratify* (i.e. split in several code sections) a NeuroLang program as deterministic and probabilistic *strata*. Stratification allows us to *programmatically* post-process and analyse results from probabilistic calculations within the same self-contained program, as illustrated in fig. 10. Moreover, this strategy can be used to supplement deterministic strata with logic programming extensions that would not necessarily be compatible with probabilistic programming [41].

PERs can either infer *marginal* or *conditional* probabilities. A marginal PER takes the form

Result(..., **PROB**, ...) :- **P1**(...) & ... & **Pn**(...)

where **P1**, ..., **Pn** are deterministic or probabilistic relational symbols, and where **PROB** is a special symbol that indicates which attribute of the resulting **Result** relation will contain the marginal probability that the rule's antecedent is true. A conditional PER is only slightly different in that it calculates the probability of a conjunction of literals being true, *given that* another conjunction of literals is true. Conditional PER take the

Figure 10: Stratification in NeuroLang. The program’s input contains both deterministic (A_1 , A_2 , and C), and probabilistic (D and E) tables.

form

```
Result(..., PROB, ...) :-
    P1(...) & ... & Pn(...) // Q1(...) & ... & Qm(...)
```

where the `//` operator applies a *probabilistic conditioning*. The `PROB` attribute of the resulting `Result` relation encodes the conditional probabilities $\mathbf{P}[P_1(\mathbf{x}) \wedge \dots \wedge P_n(\mathbf{x}) | Q_1(\mathbf{x}) \wedge \dots \wedge Q_m(\mathbf{x})]$, for each tuple $\mathbf{x}$ such that the probability is strictly positive.

NeuroLang supports explicit existential quantification of variables using rules of the form $\mathbf{P}(\mathbf{x}) :- \mathbf{Q}(\mathbf{x}) \ \& \ \text{exists}(\mathbf{y}, \mathbf{R}(\mathbf{x}, \mathbf{y}))$, where variable $\mathbf{y}$ is existentially-quantified using the language’s special symbol `exists`. The language also supports negation within its deterministic and probabilistic rules, and aggregations within its deterministic rules. An example of aggregation rule is $\mathbf{P}(\mathbf{x}, \text{count}(\mathbf{y}, \mathbf{z})) :- \mathbf{Q}(\mathbf{x}, \mathbf{y}, \mathbf{z})$, which counts, for each possible assignment of x , the number of tuples (y, z) such that $Q(x, y, z)$ is true, and stores this count in the second column of table P .

Probabilistic tables can be constructed dynamically from deterministic rules. For example, the rule $\mathbf{P}(\mathbf{x}) : \mathbf{f}(\mathbf{x}, \mathbf{y}, \mathbf{z}) :- \mathbf{Q}(\mathbf{x}, \mathbf{y}, \mathbf{z})$ constructs a probabilistic table P , assigning a probability to each tuple $(x,)$ based on function f . In the case where such rule generates multiple probabilities for the same tuple $(x,)$, an error is thrown and the user is advised to either change her probabilistic definition, or apply an aggregation function on the probabilities, such as $\text{max}(\mathbf{f}(\mathbf{x}, \mathbf{y}, \mathbf{z}))$. The antecedent of these rules must be deterministic.

Likelihood Ratio Test for NeuroLang Queries

Throughout this work, we formulate meta-analytic conditional probabilistic queries of the form $\mathbf{P}[\varphi(s) | \psi(s)]$, where $\varphi(s)$ and $\psi(s)$ are first-order logic formulas describing study-specific probabilistic events of interest; such as whether a region / network is reported by study s , or whether s is associated with a topic related to a particular psychological concept. For brevity, we write $\mathbf{P}[\varphi(s) | \psi(s)]$ instead of $\mathbf{P}[\varphi(s) = \top | \psi(s) = \top]$, where $\varphi(s)$ and $\psi(s)$ are modeled as Bernoulli random variables that have a probability of being true ($\top$) or false ($\perp$) in any possible execution of the probabilistic logic program. The formula $\psi(s)$ imposes conditions that select studies that will be included in a meta-analysis. To test the statistical dependence of $\varphi(s)$ on $\psi(s)$, we use a likelihood ratio test, whose null

(H_0) and alternative (H_1) hypotheses are

$$H_0 : \mathbf{P}[\varphi(s)|\psi(s)] = \mathbf{P}[\varphi(s)|\neg\psi(s)] = \mathbf{P}[\varphi(s)] \quad (1)$$

$$H_1 : \mathbf{P}[\varphi(s)|\psi(s)] \neq \mathbf{P}[\varphi(s)|\neg\psi(s)] \quad (2)$$

We define the likelihood ratio as $\lambda = \mathcal{L}(H_1)/\mathcal{L}(H_0)$, where $\mathcal{L}(H_1)$ and $\mathcal{L}(H_0)$ are the maximum likelihood of the observed data under the alternative and null hypotheses, defined as

$$\mathcal{L}(H_0) = \text{Bin}(k; n, \mathbf{P}[\varphi(s)])\text{Bin}(m - k; N - n, \mathbf{P}[\varphi(s)]) \quad (3)$$

$$\mathcal{L}(H_1) = \text{Bin}(k; n, \mathbf{P}[\varphi(s)|\psi(s)])\text{Bin}(m - k; N - n, \mathbf{P}[\varphi(s)|\neg\psi(s)]) \quad (4)$$

where m is the number of studies s such that $\varphi(s)$, n is the number of studies s such that $\psi(s)$, k is the number of studies s such that $\varphi(s) \wedge \psi(s)$, and N is the total number of studies within the database. As $2\log \lambda$ is asymptotically χ^2 distributed with 1 degree of freedom [57], it provides an estimate of the false-positive rate when rejecting the null hypothesis.

5 Acknowledgments

This work was funded by the ERC-2017-STG NeuroLang grant. We are grateful to Guillaume Favelier who helped with some of our visualisations. We are also grateful to Jonas Renault who worked on optimising NeuroLang’s engine and developing its web-based application. The source code of NeuroLang is openly available on GitHub at <https://github.com/NeuroLang/NeuroLang>.

6 Conflict of Interest Statement

All authors declare no conflict of interest for this manuscript.

References

- [1] Pantelis Samartsidis, Silvia Montagna, Timothy D. Johnson, and Thomas E. Nichols. The coordinate-based meta-analysis of neuroimaging data. *Statist. Sci.*, 32(4):580–599, November 2017. ISSN 0883-4237. doi: 10.1214/17-sts624. URL <https://doi.org/10.1214/17-sts624>.
- [2] Angela R. Laird, Jack L. Lancaster, and Peter T. Fox. BrainMap: The social evolution of a human brain mapping database. *NI*, 3(1):065–078, 2005. ISSN 1539-2791. doi: 10.1385/ni:3:1:065. URL <https://doi.org/10.1385/ni:3:1:065>.
- [3] Tal Yarkoni, Russell A Poldrack, Thomas E Nichols, David C Van Essen, and Tor D Wager. Large-scale automated synthesis of human functional neuroimaging data. *Nat Methods*, 8(8):665–670, June 2011. ISSN 1548-7091, 1548-7105. doi: 10.1038/nmeth.1635. URL <https://doi.org/10.1038/nmeth.1635>.
- [4] Maartje Luijten, Arnt F. Schellekens, Simone Kühn, Marise W.J. Machielse, and Guillaume Sescousse. Disruption of reward processing in addiction. *JAMA Psychiatry*, 74(4):387, April 2017. ISSN 2168-622X. doi: 10.1001/jamapsychiatry.2016.3084. URL <https://doi.org/10.1001/jamapsychiatry.2016.3084>.

- [5] Timm B. Poepl, Maximilian R. Donges, Andreas Mokros, Rainer Rupprecht, Peter T. Fox, Angela R. Laird, Danilo Bzdok, Berthold Langguth, and Simon B. Eickhoff. A view behind the mask of sanity: Meta-analysis of aberrant brain activity in psychopaths. *Mol Psychiatry*, 24(3):463–470, July 2018. ISSN 1359-4184, 1476-5578. doi: 10.1038/s41380-018-0122-5. URL <https://doi.org/10.1038/s41380-018-0122-5>.
- [6] B.T. Thomas Yeo, Fenna M. Krienen, Simon B. Eickhoff, Siti N. Yaakub, Peter T. Fox, Randy L. Buckner, Christopher L. Asplund, and Michael W.L. Chee. Functional specialization and flexibility in human association cortex. *Cereb. Cortex*, 25(10):3654–3672, September 2014. ISSN 1047-3211, 1460-2199. doi: 10.1093/cercor/bhu217. URL <https://doi.org/10.1093/cercor/bhu217>.
- [7] Matthew L. Dixon, Alejandro De La Vega, Caitlin Mills, Jessica Andrews-Hanna, R. Nathan Spreng, Michael W. Cole, and Kalina Christoff. Heterogeneity within the frontoparietal control network and its relationship to the default and dorsal attention networks. *Proc Natl Acad Sci USA*, 115(7):E1598–E1607, January 2018. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1715766115. URL <https://doi.org/10.1073/pnas.1715766115>.
- [8] Tal Yarkoni, Russell A. Poldrack, David C. Van Essen, and Tor D. Wager. Cognitive neuroscience 2.0: Building a cumulative science of human brain function. *Trends Cogn. Sci.*, 14(11):489–496, November 2010. ISSN 1364-6613. doi: 10.1016/j.tics.2010.08.004. URL <https://doi.org/10.1016/j.tics.2010.08.004>. Edition: 2010/09/29.
- [9] R. Toro, P.T. Fox, and T. Paus. Functional coactivation map of the human brain. *Cereb. Cortex*, 18(11):2553–2559, February 2008. ISSN 1047-3211, 1460-2199. doi: 10.1093/cercor/bhn014. URL <https://doi.org/10.1093/cercor/bhn014>.
- [10] Tor D. Wager, Lauren Y. Atlas, Martin A. Lindquist, Mathieu Roy, Choong-Wan Woo, and Ethan Kross. An fMRI-based neurologic signature of physical pain. *N Engl J Med*, 368(15):1388–1397, April 2013. ISSN 0028-4793, 1533-4406. doi: 10.1056/nejmoa1204471. URL <https://doi.org/10.1056/nejmoa1204471>.
- [11] Jessica R. Andrews-Hanna, Jonathan Smallwood, and R. Nathan Spreng. The default network and self-generated thought: Component processes, dynamic control, and clinical relevance. *Ann. N.Y. Acad. Sci.*, 1316(1):29–52, February 2014. ISSN 0077-8923. doi: 10.1111/nyas.12360. URL <https://doi.org/10.1111/nyas.12360>.
- [12] Peter T Fox, Jack L Lancaster, Angela R Laird, and Simon B Eickhoff. Meta-analysis in human neuroimaging: computational modeling of large-scale databases. *Annual review of neuroscience*, 37:409–434, 2014.
- [13] Nilesh Dalvi and Dan Suciu. The dichotomy of probabilistic inference for unions of conjunctive queries. *J. ACM*, 59(6):1–87, December 2012. ISSN 0004-5411, 1557-735X. doi: 10.1145/2395116.2395119. URL <https://doi.org/10.1145/2395116.2395119>. Number: 6.
- [14] Guy Van den Broeck and Dan Suciu. Query processing on probabilistic data: A survey. *FNT in Databases*, 7(3-4):197–341, 2017. ISSN 1931-7883, 1931-7891. doi: 10.1561/19000000052. URL <https://doi.org/10.1561/19000000052>. Number: 3-4.

- [15] Hans Kamp and Uwe Reyle. *From Discourse to Logic*, volume 42. Springer Netherlands, 1993. ISBN 9789401049160, 9789401120661. doi: 10.1007/978-94-011-2066-1. URL <https://doi.org/10.1007/978-94-011-2066-1>.
- [16] Tor D. Wager, Martin Lindquist, and Lauren Kaplan. Meta-analysis of functional neuroimaging data: Current and future directions. *Soc. Cogn. Affect. Neur.*, 2(2): 150–158, June 2007. ISSN 1749-5024, 1749-5016. doi: 10.1093/scan/nsm015. URL <https://doi.org/10.1093/scan/nsm015>.
- [17] Simon B. Eickhoff, Danilo Bzdok, Angela R. Laird, Florian Kurth, and Peter T. Fox. Activation likelihood estimation meta-analysis revisited. *Neuroimage*, 59(3): 2349–2361, February 2012. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2011.09.017. URL <https://doi.org/10.1016/j.neuroimage.2011.09.017>.
- [18] Russell A. Poldrack, Jeanette A. Mumford, Tom Schonberg, Donald Kalar, Bishal Barman, and Tal Yarkoni. Discovering relations between mind, brain, and mental disorders using topic mapping. *PLoS Comput Biol*, 8(10):e1002707, October 2012. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1002707. URL <https://doi.org/10.1371/journal.pcbi.1002707>.
- [19] Valentin Iovene, Gaston E Zanitti, and Demian Wassermann. Complex coordinate-based meta-analysis with probabilistic programming. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(1):223–231, May 2021. URL <https://ojs.aaai.org/index.php/AAAI/article/view/16096>.
- [20] Kamalaker Dadi, Gaël Varoquaux, Antonia Machlouzarides-Shalit, Krzysztof J. Gorgolewski, Demian Wassermann, Bertrand Thirion, and Arthur Mensch. Fine-grain atlases of functional modes for fMRI analysis. *Neuroimage*, 221:117126, November 2020. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2020.117126. URL <https://doi.org/10.1016/j.neuroimage.2020.117126>.
- [21] B.T. Thomas Yeo, Fenna M. Krienen, Jorge Sepulcre, Mert R. Sabuncu, Danial Lashkari, Marisa Hollinshead, Joshua L. Roffman, Jordan W. Smoller, Lilla Zöllei, Jonathan R. Polimeni, Bruce Fischl, Hesheng Liu, and Randy L. Buckner. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J. Neurophysiol.*, 106(3):1125–1165, September 2011. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.00338.2011. URL <https://doi.org/10.1152/jn.00338.2011>. Publisher: American Physiological Society Bethesda, MD.
- [22] R. Nathan Spreng, Jorge Sepulcre, Gary R. Turner, W. Dale Stevens, and Daniel L. Schacter. Intrinsic architecture underlying the relations among the default, dorsal attention, and frontoparietal control networks of the human brain. *J. Cognitive Neurosci.*, 25(1):74–86, January 2013. ISSN 0898-929X, 1530-8898. doi: 10.1162/jocn.a_00281. URL https://doi.org/10.1162/jocn.a_00281.
- [23] Radek Ptak, Armin Schnider, and Julia Fellrath. The dorsal frontoparietal network: A core system for emulated action. *Trends Cogn. Sci.*, 21(8):589–599, August 2017. ISSN 1364-6613. doi: 10.1016/j.tics.2017.05.002. URL <https://doi.org/10.1016/j.tics.2017.05.002>.

- [24] Lang Chen, Demian Wassermann, Daniel A. Abrams, John Kochalka, Guillermo Gallardo-Diez, and Vinod Menon. The visual word form area (VWFA) is part of both language and attention circuitry. *Nat Commun*, 10(1):5601, December 2019. ISSN 2041-1723. doi: 10.1038/s41467-019-13634-z. URL <https://doi.org/10.1038/s41467-019-13634-z>.
- [25] Maurizio Corbetta, Gaurav Patel, and Gordon L. Shulman. The reorienting system of the human brain: From environment to theory of mind. *Neuron*, 58(3):306–324, May 2008. ISSN 0896-6273. doi: 10.1016/j.neuron.2008.04.017. URL <https://doi.org/10.1016/j.neuron.2008.04.017>.
- [26] John Duncan. The multiple-demand (MD) system of the primate brain: Mental programs for intelligent behaviour. *Trends Cogn. Sci.*, 14(4):172–179, April 2010. ISSN 1364-6613. doi: 10.1016/j.tics.2010.01.004. URL <https://doi.org/10.1016/j.tics.2010.01.004>.
- [27] Carrisa V Cocuzza, Takuya Ito, Douglas Schultz, Danielle S Bassett, and Michael W Cole. Flexible coordinator and switcher hubs for adaptive task control. *Journal of Neuroscience*, 40(36):6949–6968, 2020.
- [28] Tal Yarkoni, Jeremy R. Gray, Elizabeth R. Chrsatil, Deanna M. Barch, Leonard Green, and Todd S. Braver. Sustained neural activity associated with cognitive control during temporally extended decision making. *Cognitive Brain Res.*, 23(1): 71–84, April 2005. ISSN 0926-6410. doi: 10.1016/j.cogbrainres.2005.01.013. URL <https://doi.org/10.1016/j.cogbrainres.2005.01.013>.
- [29] Tara A. Niendam, Angela R. Laird, Kimberly L. Ray, Y. Monica Dean, David C. Glahn, and Cameron S. Carter. Meta-analytic evidence for a superordinate cognitive control network subserving diverse executive functions. *Cogn Affect Behav Neurosci*, 12(2):241–268, January 2012. ISSN 1530-7026, 1531-135X. doi: 10.3758/s13415-011-0083-5. URL <https://doi.org/10.3758/s13415-011-0083-5>.
- [30] Jérôme Dockes. *Statistical models for comprehensive meta-analysis of neuroimaging studies*. Theses, Université Paris-Saclay, November 2019. Issue: 2019SACLT048.
- [31] Vinod Menon, Guillermo Gallardo, Mark A Pinsk, Van-Dang Nguyen, Jing-Rebecca Li, Weidong Cai, and Demian Wassermann. Microstructural organization of human insula is linked to its macrofunctional circuitry and predicts cognitive control. *eLife*, 9:e53470, June 2020. ISSN 2050-084X. doi: 10.7554/elife.53470. URL <https://doi.org/10.7554/elife.53470>.
- [32] R Nathan Spreng, W Dale Stevens, Jon P Chamberlain, Adrian W Gilmore, and Daniel L Schacter. Default network activity, coupled with the frontoparietal control network, supports goal-directed cognition. *Neuroimage*, 53(1):303–317, 2010.
- [33] Danilo Bzdok, Gaël Varoquaux, Olivier Grisel, Michael Eickenberg, Cyril Poupon, and Bertrand Thirion. Formal models of the network co-occurrence underlying mental operations. *PLoS Comput Biol*, 12(6):e1004994, June 2016. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1004994. URL <https://doi.org/10.1371/journal.pcbi.1004994>.

- [34] Lucina Q Uddin, BT Thomas Yeo, and R Nathan Spreng. Towards a universal taxonomy of macro-scale functional human brain networks. *Brain topography*, 32(6): 926–942, 2019.
- [35] Marcus E. Raichle. The brain’s default mode network. *Annu. Rev. Neurosci.*, 38(1):433–447, July 2015. ISSN 0147-006X, 1545-4126. doi: 10.1146/annurev-neuro-071013-014030. URL <https://doi.org/10.1146/annurev-neuro-071013-014030>.
- [36] Simone Vossel, Joy J. Geng, and Gereon R. Fink. Dorsal and ventral attention systems. *Neuroscientist*, 20(2):150–159, July 2013. ISSN 1073-8584, 1089-4098. doi: 10.1177/1073858413494269. URL <https://doi.org/10.1177/1073858413494269>.
- [37] Alecia C Vogel, Fran M Miezin, Steven E Petersen, and Bradley L Schlaggar. The putative visual word form area is functionally connected to the dorsal attention network. *Cerebral cortex*, 22(3):537–549, 2012.
- [38] Kevin S. Weiner and Karl Zilles. The anatomical and functional specialization of the fusiform gyrus. *Neuropsychologia*, 83:48–62, March 2016. ISSN 0028-3932. doi: 10.1016/j.neuropsychologia.2015.06.033. URL <https://doi.org/10.1016/j.neuropsychologia.2015.06.033>.
- [39] Fenna M. Krienen, B.T. Thomas Yeo, and Randy L. Buckner. Reconfigurable task-dependent functional coupling modes cluster around a core functional architecture. *Phil. Trans. R. Soc. B*, 369(1653):20130526, October 2014. ISSN 0962-8436, 1471-2970. doi: 10.1098/rstb.2013.0526. URL <https://doi.org/10.1098/rstb.2013.0526>.
- [40] Russell A Poldrack, Aniket Kittur, Donald Kalar, Eric Miller, Christian Seppa, Yolanda Gil, D Stott Parker, Fred W Sabb, and Robert M Bilder. The cognitive atlas: toward a knowledge foundation for cognitive neuroscience. *Frontiers in neuroinformatics*, 5:17, 2011.
- [41] Gaston E Zanitti, Yamil Soto, Valentin Iovene, Maria Vanina Martinez, Ricardo O Rodriguez, Gerardo I Simari, and Demian Wassermann. Scalable Query Answering under Uncertainty to Neuroscientific Ontological Knowledge: The NeuroLang Approach. working paper or preprint, April 2021. URL <https://hal.inria.fr/hal-03187887>.
- [42] K.J. Friston, L. Harrison, and W. Penny. Dynamic causal modelling. *Neuroimage*, 19(4):1273–1302, August 2003. ISSN 1053-8119. doi: 10.1016/s1053-8119(03)00202-7. URL [https://doi.org/10.1016/s1053-8119\(03\)00202-7](https://doi.org/10.1016/s1053-8119(03)00202-7).
- [43] Rotem Botvinik-Nezer, Felix Holzmeister, Colin F. Camerer, Anna Dreber, Juer-gen Huber, Magnus Johannesson, Michael Kirchler, Roni Iwanir, Jeanette A. Mumford, R. Alison Adcock, Paolo Avesani, Blazej M. Baczkowski, Aahana Bajracharya, Leah Bakst, Sheryl Ball, Marco Barilari, Nadège Bault, Derek Beaton, Julia Beitner, Roland G. Benoit, Ruud M.W.J. Berkers, Jamil P. Bhanji, Bharat B. Biswal, Sebastian Bobadilla-Suarez, Tiago Bortolini, Katherine L. Bottenhorn, Alexander Bowring, Senne Braem, Hayley R. Brooks, Emily G. Brudner, Cristian B. Calderon, Julia A. Camilleri, Jaime J. Castrellon, Luca Cecchetti, Edna C. Cieslik, Zachary J.

Cole, Olivier Collignon, Robert W. Cox, William A. Cunningham, Stefan Czoschke, Kamalaker Dadi, Charles P. Davis, Alberto De Luca, Mauricio R. Delgado, Lysia Demetriou, Jeffrey B. Dennison, Xin Di, Erin W. Dickie, Ekaterina Dobryakova, Claire L. Donnat, Juergen Dukart, Niall W. Duncan, Joke Durnez, Amr Eed, Simon B. Eickhoff, Andrew Erhart, Laura Fontanesi, G. Matthew Fricke, Shiguang Fu, Adriana Galván, Remi Gau, Sarah Genon, Tristan Glatard, Enrico Glerean, Jelle J. Goeman, Sergej A.E. Golowin, Carlos González-García, Krzysztof J. Gorgolewski, Cheryl L. Grady, Mikella A. Green, João F. Guassi Moreira, Olivia Guest, Shabnam Hakimi, J. Paul Hamilton, Roeland Hancock, Giacomo Handjaras, Bronson B. Harry, Colin Hawco, Peer Herholz, Gabrielle Herman, Stephan Heunis, Felix Hoffstaedter, Jeremy Hogeveen, Susan Holmes, Chuan-Peng Hu, Scott A. Huettel, Matthew E. Hughes, Vittorio Iacovella, Alexandru D. Iordan, Peder M. Isager, Ayse I. Isik, Andrew Jahn, Matthew R. Johnson, Tom Johnstone, Michael J.E. Joseph, Anthony C. Juliano, Joseph W. Kable, Michalis Kassinosopoulos, Cemal Koba, Xiang-Zhen Kong, Timothy R. Kosciuk, Nuri Erkut Kucukboyaci, Brice A. Kuhl, Sebastian Kupek, Angela R. Laird, Claus Lamm, Robert Langner, Nina Lauharatanahirun, Hongmi Lee, Sangil Lee, Alexander Leemans, Andrea Leo, Elise Lesage, Flora Li, Monica Y.C. Li, Phui Cheng Lim, Evan N. Lintz, Schuyler W. Liphardt, Annabel B. Losecaat Vermeer, Bradley C. Love, Michael L. Mack, Norberto Malpica, Theo Marins, Camille Maumet, Kelsey McDonald, Joseph T. McGuire, Helena Melero, Adriana S. Méndez Leal, Benjamin Meyer, Kristin N. Meyer, Glad Mihai, Georgios D. Mitsis, Jorge Moll, Dylan M. Nielson, Gustav Nilsson, Michael P. Notter, Emanuele Olivetti, Adrian I. Onicas, Paolo Papale, Kaustubh R. Patil, Jonathan E. Peelle, Alexandre Pérez, Doris Pischella, Jean-Baptiste Poline, Yanina Prystauka, Shruti Ray, Patricia A. Reuter-Lorenz, Richard C. Reynolds, Emiliano Ricciardi, Jenny R. Rieck, Anais M. Rodriguez-Thompson, Anthony Romy, Taylor Salo, Gregory R. Samanez-Larkin, Emilio Sanz-Morales, Margaret L. Schlichting, Douglas H. Schultz, Qiang Shen, Margaret A. Sheridan, Jennifer A. Silvers, Kenny Skagerlund, Alec Smith, David V. Smith, Peter Sokol-Hessner, Simon R. Steinkamp, Sarah M. Tashjian, Bertrand Thirion, John N. Thorp, Gustav Tinghög, Loreen Tisdall, Steven H. Tompson, Claudio Toro-Serey, Juan Jesus Torre Tresols, Leonardo Tozzi, Vuong Truong, Luca Turella, Anna E. van 't Veer, Tom Verguts, Jean M. Vettel, Sagana Vijayarajah, Khoi Vo, Matthew B. Wall, Wouter D. Weeda, Susanne Weis, David J. White, David Wisniewski, Alba Xifra-Porxas, Emily A. Yearling, Sangsuk Yoon, Rui Yuan, Kenneth S.L. Yuen, Lei Zhang, Xu Zhang, Joshua E. Zosky, Thomas E. Nichols, Russell A. Poldrack, and Tom Schonberg. Variability in the analysis of a single neuroimaging dataset by many teams. *Nature*, 582 (7810):84–88, June 2020. ISSN 1476-4687. doi: 10.1038/s41586-020-2314-9. URL <https://doi.org/10.1038/s41586-020-2314-9>.

- [44] Serge Abiteboul and Richard Hull. Data functions, datalog and negation. In *Proceedings of the 1988 ACM SIGMOD international conference on Management of data - SIGMOD '88*, SIGMOD '88, pages 143–153, New York, NY, USA, 1988. ACM Press. ISBN 0897912683. doi: 10.1145/50202.50218. URL <https://doi.org/10.1145/50202.50218>.
- [45] Robert Fink and Dan Olteanu. Dichotomies for queries with negation in probabilistic databases. *ACM Trans. Database Syst.*, 41(1):1–47, April 2016. ISSN 0362-5915, 1557-4644. doi: 10.1145/2877203. URL <https://doi.org/10.1145/2877203>.

- [46] Moshe Y. Vardi. The complexity of relational query languages (Extended abstract). In *Proceedings of the fourteenth annual ACM symposium on Theory of computing - STOC '82*, pages 137–146, San Francisco, California, United States, 1982. ACM Press. ISBN 0897910702. doi: 10.1145/800070.802186. URL <https://doi.org/10.1145/800070.802186>.
- [47] Adnan Darwiche, Pierre Marquis, Dan Suciu, and Stefan Szeider. Recent trends in knowledge compilation (dagstuhl seminar 17381). In *Dagstuhl Reports*, volume 7. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- [48] Veronika I. Müller, Edna C. Cieslik, Angela R. Laird, Peter T. Fox, Joaquim Radua, David Mataix-Cols, Christopher R. Tench, Tal Yarkoni, Thomas E. Nichols, Peter E. Turkeltaub, Tor D. Wager, and Simon B. Eickhoff. Ten simple rules for neuroimaging meta-analysis. *Neuroscience & Biobehavioral Reviews*, 84:151–161, January 2018. ISSN 0149-7634. doi: 10.1016/j.neubiorev.2017.11.012. URL <https://doi.org/10.1016/j.neubiorev.2017.11.012>.
- [49] S. Abiteboul, Richard Hull, and Victor Vianu. *Foundations of databases*. Addison-Wesley, Reading, Mass, 1995. ISBN 978-0-201-53771-0.
- [50] Anton Dries, Angelika Kimmig, Wannes Meert, Joris Renkens, Guy Van den Broeck, Jonas Vlasselaer, and Luc De Raedt. ProbLog2: Probabilistic logic programming. In *Machine Learning and Knowledge Discovery in Databases*, volume 9286, pages 312–315. Springer International Publishing, Cham, 2015. ISBN 9783319234601, 9783319234618. doi: 10.1007/978-3-319-23461-8_37. URL https://doi.org/10.1007/978-3-319-23461-8_37. Series Title: Lecture Notes in Computer Science.
- [51] Joost Vennekens, Marc Denecker, and Maurice Bruynooghe. CP-logic: A language of causal probabilistic events and its relation to logic programming. *Theor. Pract. Log. Prog.*, 9(3):245–308, May 2009. ISSN 1471-0684, 1475-3081. doi: 10.1017/s1471068409003767. URL <https://doi.org/10.1017/s1471068409003767>.
- [52] Taisuke Sato. A statistical learning method for logic programs with distribution semantics. In *Proceedings of the 12th international conference on logic programming (ICLP'95)*, pages 715–729. MIT Press, 1995.
- [53] Luc De Raedt and Angelika Kimmig. Probabilistic (logic) programming concepts. *Mach Learn*, 100(1):5–47, May 2015. ISSN 0885-6125, 1573-0565. doi: 10.1007/s10994-015-5494-z. URL <https://doi.org/10.1007/s10994-015-5494-z>.
- [54] Tian Sang, Paul Beame, and Henry A. Kautz. Performing Bayesian inference by weighted model counting. In *AAAI*, pages 475–482, 2005. URL <http://www.aaai.org/Library/AAAI/2005/aaai05-075.php>.
- [55] Adnan Darwiche. SDD: A new canonical representation of propositional knowledge bases. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence - Volume Volume Two, IJCAI'11*, pages 819–826. AAAI Press, 2011.
- [56] Dan Suciu, Dan Olteanu, Christopher Ré, and Christoph Koch. Probabilistic databases. *Synthesis Lectures on Data Management*, 3(2):1–180, May 2011. ISSN

2153-5418, 2153-5426. doi: 10.2200/s00362ed1v01y201105dtm016. URL <https://doi.org/10.2200/s00362ed1v01y201105dtm016>.

- [57] Alexander McFarlane Mood. *Introduction to the Theory of Statistics*. McGraw-hill, 1950.