

The Diagrammatic AI Language (DIAL): Version 0.1

Guy Marshall and André Freitas

School of Computer Science
University of Manchester
United Kingdom

December 2018

Abstract

Currently, there is no consistent model for visually or formally representing the architecture of AI systems. This lack of representation brings interpretability, correctness and completeness challenges in the description of existing models and systems. DIAL (The Diagrammatic AI Language) has been created with the aspiration of being an “engineering schematic” for AI Systems. It is presented here as a starting point for a community dialogue towards a common diagrammatic language for AI Systems.

1 Introduction

1.1 Motivation

The development of real-world Artificial Intelligence systems requires the complex coordination of multiple components. Currently, there is no consistent model for visually or formally representing the architecture of AI systems. This lack of representation brings interpretability, correctness and completeness challenges in the description of existing models and systems.

In the context of scientific communication, most approaches and systems today are described by a combination of arbitrary diagrammatic elements, algorithms, formulae and natural language descriptions. In this paradigm, there is little consistency on abstraction levels (for example, it is common for diagrams to refer to formulae terms) or notation.

From the perspective of scientific practice, these limitations challenge dialogue, transparency and reproducibility. This work outlines the initial specification, design and application of a diagrammatic language for describing complex, multi-component AI systems. The goal of the language, named ‘DIAL’ (Diagrammatic AI Language), is to reduce the communication effort between AI

researchers and engineers for describing, interpreting and reasoning over complex AI systems.

1.2 Background

The idea of applying diagrammatic representations to AI is not new. In 1985, Aaron Sloman writes about the need for representational formalisms in AI, praising diagrams for controlling the search space [15]. In a related way, the Sapir-Wharf hypothesis states that the structure of native language determines the native speaker's perception and categorisation experience, a conjecture which could be extended to diagrammatic languages. Following this thinking, we argue that for AI researchers and engineers, diagrammatic structures can determine perception and categorisation of AI systems and subsystems. As such, a diagrammatic language (such as DIAL) may aid in reasoning about AI systems.

In the remainder of Section 1 we will look at the main structured alternative diagrammatic approach (UML) and show the current heterogeneity issue with AI System diagrams. This will provide the necessary grounding to propose requirements for AI System diagrams, and introduce DIAL as a potential solution to these issues. The main diagrammatic components of DIAL are then described, formalised and exemplified. As this document focuses on the specification of DIAL v. 0.1, we concentrate on the narrative around the requirements and design for DIAL, seeking the engagement of a community around DIAL and the collection of critiques and contributions.

1.3 Key Benefits of DIAL

The core benefits aimed for DIAL are:

1. Having a representation which supports AI systems design, reasoning and learning.
2. Increasing the consistency and efficiency of the communication around AI Systems.

The rapid evolution of techniques have been defining an increasingly dynamic and fragmented AI landscape. The plethora of proposed models and tools and their associated complexity brings an increasing difficulty to interpret, contrast, compare and select models in a principled and critical manner. DIAL is designed using a combined bottom-up/top-down design approach. In the bottom-up component, reoccurring features observed in existing diagrams are amalgamated into the representation, complemented by the translation of elements represented as text into the integrated diagrammatic representation. In the top-down component, high-level requirements are introduced, focusing on the completeness and interpretability aspects of the representation.

DIAL aims at providing the means for a self-contained description of the core functional elements of an AI system from both research and engineering

perspectives. Supporting datasets, data management infrastructures, gold standards and evaluation metrics are first-class citizens of the representation. Additionally, DIAL aims by design to provide a comprehensive set of components within the full AI spectrum, not privileging a particular category of techniques (e.g. Neural Networks).

The self-contained and concise aspects of DIAL can be used as a foundation for a possible way to materialise *nano-publications* in the context of AI. Using DIAL, readers can navigate through a concise representation of an AI system or experiment, freeing the AI practitioner from the need to go through heterogeneous diagrammatic and textual descriptions. This format supports the ability to quickly understand novel research and approaches, and be refreshed of them quickly without having to reread a supporting textual description. In the short-term, it is hoped that the usage of DIAL will facilitate more effective consumption of AI System concepts.

1.4 Why UML isn't a Good Answer

Within software more broadly, the communication and representation of architectures is supported [14] with standards such as UML [3] and DARWIN [8].

For more specific applications, different diagrammatic representations are also being created. For “user requirements specification”, nine deficiencies in UML have been identified, including that “UML cannot model the behavior of high-level system components...” [5]. In dynamic modeling, Otero [11] found Open Modeling Language worked better than UML for both quality and comprehension ratings.

UML is also seen as overly complex, with UML Class Diagrams having a graphic complexity of 14 (where the human cognitive limit is six) [9]. Perhaps the most compelling argument against the suitability of UML comes from a pragmatic space. The existence of UML is quite common knowledge, and yet is not being used by AI practitioners, perhaps for reasons described in the above paragraph.

Within AI, the increasing complexity of systems increases the associated effort in communicating and understanding the fundamental functional components of a system. In the context of AI research and practice, there is a natural entanglement between the scientific and engineering discourses. Additionally, the design of an AI system will involve the consideration of multiple abstraction levels, ranging from a formal mathematical description level, to their algorithmic realisation and upstream software materialisation.

While this task imposes major challenges in a pure ML pipeline, with architectural elements, hyper-parameters, feature engineering and associated datasets, this effort grows exponentially in systems which involve the coordination of multiple and heterogeneous classification, analytic and inference steps. In this context, a diagrammatic representation for AI can provide an interpretation mechanism which captures the core methods, resources and interdependencies employed within the AI task in a holistic manner.

1.5 The Problem of Diagrammatic Heterogeneity

There is a large variety of AI Diagrams being used within the AI community at present. Taking the sub-area of Neural Networks as an example, a large variety of neural architectures have emerged recently. Some initiatives to collect and describe them systematically have surfaced <http://www.asimovinstitute.org/neural-network-zoo/>. However, as the author states: “One problem with drawing them [Neural Networks] as node maps: it doesn’t really show how they’re used”. Tensorflow has graph visualisations (see TensorBoard) which are also instructive, at a lower level of abstraction, and tied to a specific technology.

In order to quantify this existing heterogeneity, we performed a short systematic literature analysis on the use of diagrams to represent NNs in the context of NLP systems.

The analysis started by selecting 30 randomly sample papers from the overall 331 accepted papers from COLING 2018. For the sampled papers, diagrams of Neural Networks were manually extracted, and the diagrams categorised. Only the diagrams and the text or formulae included within those diagrams is analysed. The supporting context is purposefully ignored in order to examine the effectiveness of the diagrammatic representation itself. COLING 2018 was chosen exclusively in order to ensure the same social context. Prior to the study, we identified a set of core NN concepts often articulated through diagrammatic means, which have been captured in Table 1.

Of the 30 papers selected at random, 23 (77%) were manually classified as being supported by a NN-based classifier. Note that this was not filtered for Neural Network papers, and this perhaps serves to highlight the prevalence of this technique at this time. Of the papers in some way concerned with Neural Networks, 78% include a diagram of a Neural Network. Table 1 shows the distribution of Neural Network diagrams with each conceptual property.

Table 1: Neural Network Diagram Summary

CONCEPT OR ATTRIBUTE	FREQUENCY
ABSTRACTION: HIGH LEVEL SYSTEM	67%
ABSTRACTION: INTERMEDIATE LEVEL	50%
ABSTRACTION: COMPONENT LEVEL	28%
ABSTRACTION: MIXED LEVELS	6%
MULTIPLE DIAGRAMS	44%
PROCESSING LAYER ORGANISATION	61%
CLASSIFIER TYPE	33%
EQUATIONAL OR MATHEMATICAL ELEMENTS	72%
VECTOR OR TENSOR OPERATION	33%
EXAMPLE INPUT, OUTPUT OR PROCESSING	39%
DIMENSIONALITY	11%
LANGUAGE ELEMENTS	94%
RECURRENT OPERATIONS	94%

The vast majority of diagrams are high level diagrams, complemented by textual labels (e.g. LSTM). However, given the diagrams from a paper at random, it is often impossible to deduce fundamental aspects of the system, such as the functional elements. Table 1 shows that few diagrams include classifier types, nor dimensional or compositional information.

DLPaper2Code categorises Deep Learning model diagrams into five categories: Neuron plot, 2D box, Stacked2D box, 3D box and Pipeline plot [13]. The authors showed this coarse-grained classification useful, though this obfuscates the cognitive complexity involved with the variety of representations within each category. There is a huge variety of semiotic principles utilised in different ways, from semiotic choices such as how colour is used, through to “icon” choices, and even in terms of semantic content (i.e. what they choose to represent).

An effective way to informally see the heterogeneity issue is by observing concrete examples of diagrams in use. A simple randomised selection of papers from the main AI conferences can demonstrate this. In the following examination, we include a brief commentary on some diagrams presented as part of papers at EMNLP 2017. Figures 1, 2 and 3 depict examples of different diagrammatic design choices, including different abstraction levels, content, and representation.

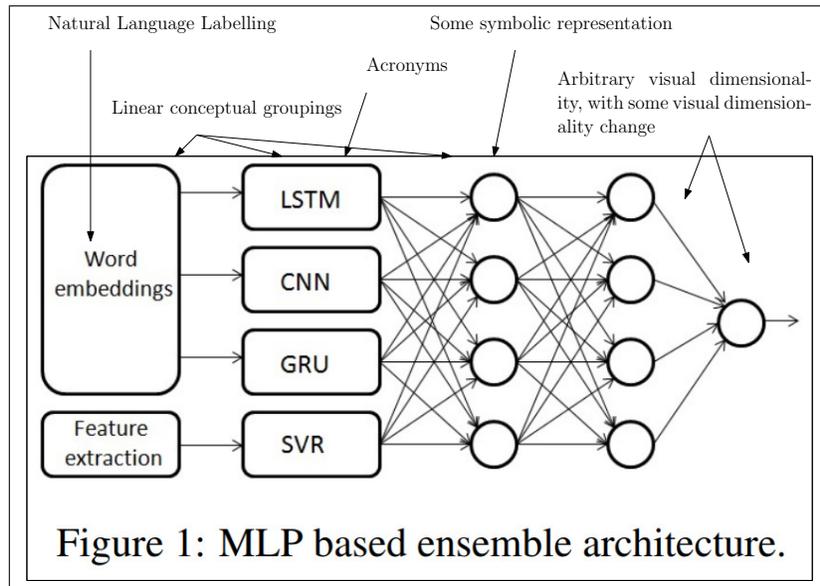


Figure 1: Diagram for a Multilayer Perceptron Ensemble [2]

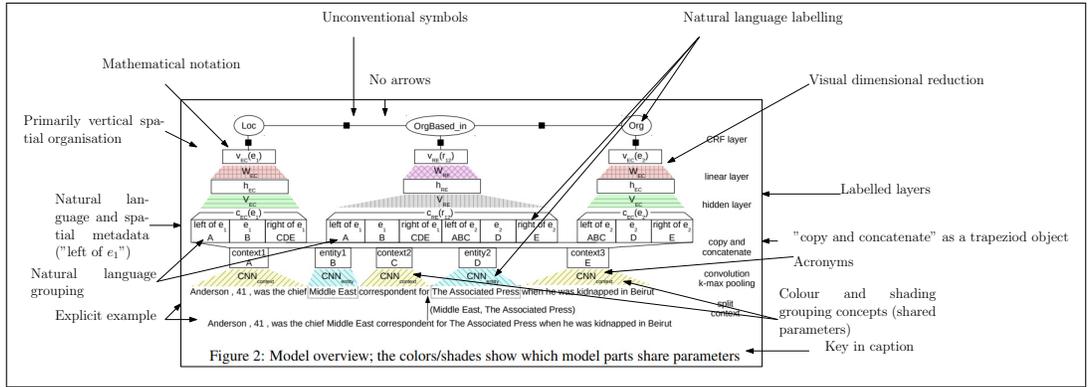


Figure 2: Global Normalization of CNN [1]

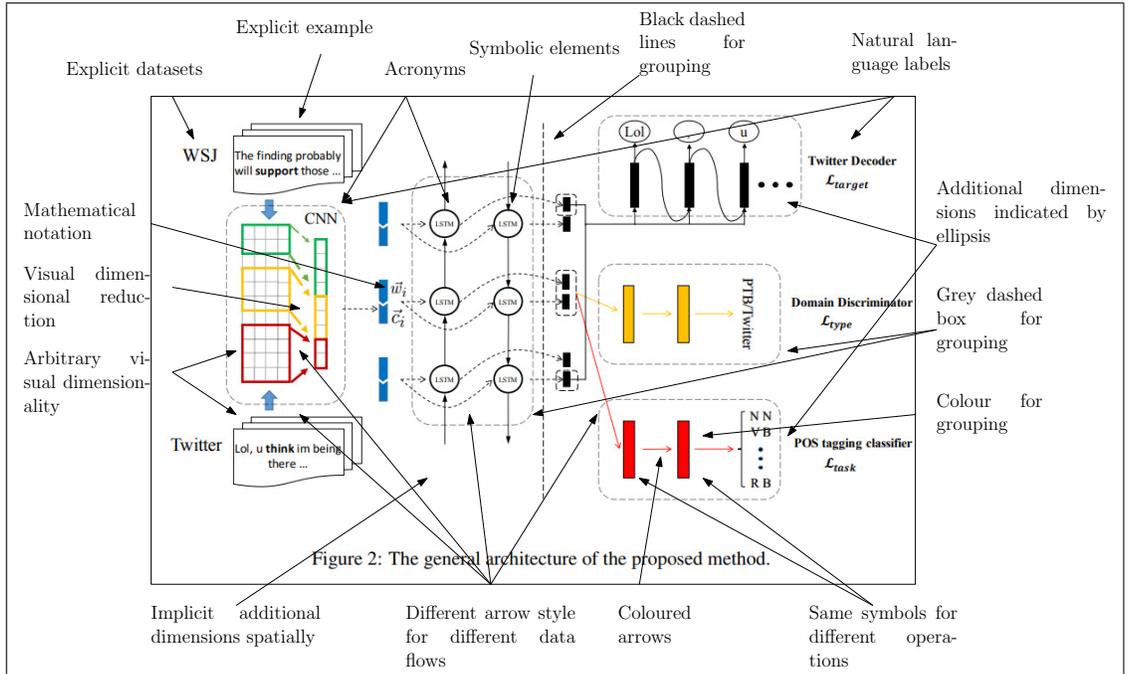


Figure 3: Part-of-Speech tagging with adversarial neural networks [7]

2 Requirements

AI systems orchestrate design choices of *semantic representation*, *algorithms*, *software architecture* and *data resources*. As such, there are multiple levels of abstraction and different use cases to be considered as part of understanding the requirements for a diagrammatic representation. Rather than engage directly with users in the initial stages, we have taken the opportunity to seed this conversation with a formal semiotic slant which will allow for additional thought, formalisation and rigor as the diagrammatic representation evolves.

2.1 Use Cases and Target Questions

In our consideration of the utility of diagrammatic representations for AI systems we adopt two basic personas: That of an *AI Researcher* attempting to review and understand the contribution of the system, and that of an *AI Systems Engineer*, attempting to understand, reproduce or adapt the system architecture. The following list materialise some of the questions which the diagrammatic representation should answer directly.

1. What are the main semantic representation choices in the system? (Relating to Category Theory's *objects* [4])
 - (a) Latent vs explicit representations, embeddings, data models, compositional models, classification categories, feature types, data integration points.
 - (b) How the data representation changes along the system?
2. Which are the main AI algorithmic and functional choices, themes and patterns in the system? (Relating to Category Theory's *morphisms*)
 - (a) Which recurring classification categories are used and what is their underlying architecture?
 - (b) What are the functional composition patterns and their interdependencies?
 - (c) What are the hyperparameters configurations, and the optimisation ranges?
3. What are the main software engineering choices/themes/patterns in the system?
 - (a) What are the temporal dependencies of the components (e.g. asynchronous vs synchronous) and the emerging main architectural themes (e.g. event-based, workflow)?
 - (b) How are functional units deployed (e.g. in memory, services, or microservices containerised) and how do they communicate?
 - (c) What are the human interaction points in the system?

4. What are the main data management choices/themes/patterns (databases/data model, stream/batch processing) in the system?
5. Which performance mechanisms are employed to ensure an efficient data access (query processing, indexing strategies)?
6. What is the quality delivered by each component and under which evaluation metric?
7. What and how are the data resources (knowledge bases, linguistic resources and gold standards) used in the system?

2.2 DIAL Requirements

The target questions lead to the following requirements for the representation:

1. To target an abstraction-level which is optimal for complex and multi-component AI systems.
2. To have a primary abstraction level of functional components and data transformation across the system.
3. To keep a coherent perspective of the whole.
4. To focus beyond the description of Machine Learning techniques.
5. To utilise non-linguistic symbolic language for recurring primitives (minimization of verbal elements).
6. To be extensible to relevant software architecture and data management aspects.
7. To communicate the performance of each component.
8. To be cognitively efficient rather than informationally dense.

Quantification, measurement and analysis of the above is intended to take place as DIAL evolves. At this stage, DIAL attempts to satisfy the “representational requirements” but not to completely address all of the “target questions”, see Section 3.

3 Roadmap

In order to allow for different concepts and levels of abstraction, whilst remaining compatible, we are presenting DIAL as a set of dialects. The dialects presented here are first iterations of:

- DIAL-SYS, providing the core common language to describe AI Systems at a high level.

- DIAL-NN, extending to Neural Network components.

The scope of DIAL-SYS is to describe complex (multi-component) AI systems in operation. DIAL-NN extends the DIAL-SYS *symbolset* to facilitate diagrammatic representation of neural network aspects of AI Systems. This area was chosen to illustrate the possibility of dialects, and due to the prevalence of neural networks architectures in current literature. Additional dialects to cover other aspects of AI Systems, such as logic will be published soon. To re-iterate, DIAL, encompassing these dialects, is not presented as a complete framework, but rather a starting point for community dialogue and collaborative construction. It is expected that significant work will be required in order to make DIAL-SYS and DIAL-NN optimal for the AI Community. Once this is in a reasonable shape, and if it proves useful, additional dialects would be expected to evolve in order to tackle different aspects of AI Systems. Future dialects will include:

- DIAL-DB, focused on data analysis and data management.
- DIAL-LOG, extending to logic formalisms and systems.
- DIAL-ML, specialised for Machine-Learning
- DIAL-PGM, for probabilistic graph models such as Bayesian Networks
- DIAL-BIZ, to extend to business or application contexts
- DIAL-SEM, for semantic knowledge representations [12]

DIAL-SYS will be developed to incorporate temporal elements, such as execution performance. Extension of DIAL to include representation for “Human-in-the-loop”, including external interaction points such as crowd-sourcing for training, will also be required as DIAL matures.

Figure 4 contains an incomplete view of some of the expected DIAL dialects (DIAL-ects). Dark circles show existing work, and lines between the circles indicate significant dependency between dialects. The prioritisation and timeline will be based on the comments received, and led by the future stewards of this standardisation effort who it is hoped will come forward (see Section 9).

4 Applications

With a more defined scope for DIAL, we continue to detail some of the anticipated applications of DIAL. The following subsections explore some of the possibilities of DIAL to support AI community members.

4.1 Nano-publications

With such a high pace of change and volume of conference papers in AI, it is extremely difficult to keep on top of the latest developments in the field. One

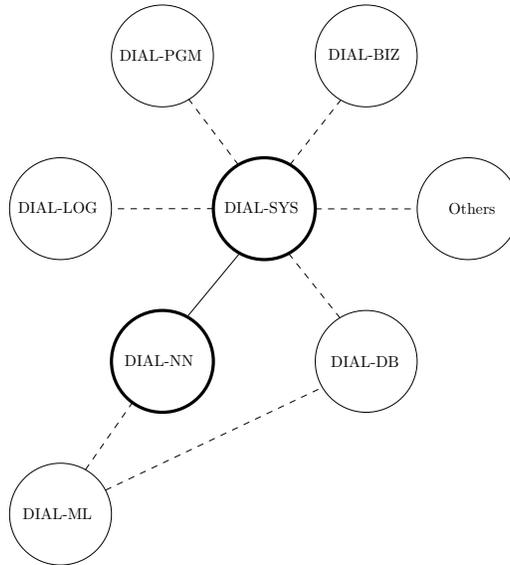


Figure 4: DIAL Architectural Map

nano-publication, consisting of a DIAL diagram, can summarise the composition of the functional elements of the proposed model, their associated performance and core architectural decisions as well as the contribution.

4.2 Software specification

DIAL, as a language, allows for specification of software by prioritising recurring AI components in the representation. This focus on AI by design can provide a more concise representation, focusing on the articulation of the AI-related design decisions (in contrast with UML).

4.3 Designing AI Systems

There is the potential, with substantial development effort, to extend DIAL to be able to simulate AI Systems in a way analogous to Electronic Design Automation for chips. This could allow for rapid prototyping and facilitate reasoning about novel architectures. Differently from TensorBoard, this would be created to support multiple different dataflow libraries, and aim to be technology agnostic.

4.4 Estimating the Impact of AI Systems in Businesses

DIAL can serve as a lower entry barrier for communicating AI functionalities and their associated performance levels to non-AI experts stakeholders. This can serve as a mechanism to communicate requirements and to synchronize expectations between business and technical profiles. This expected that DIAL

can improve the formalisation of the requirements (with a lower effort) and the quality of AI software estimation.

4.5 Education

In education, the creation of curricula describing the different perspectives around the construction of AI systems is still an open challenge. DIAL can reduce the barriers for conceptualising AI systems under a complex/multi-technique perspective, providing an integrative view of AI. Moreover, currently there is a substantial gap between the discussions within the research community and their industrial application. By reducing the interpretation bottleneck of existing scientific contributions, it is expected that DIAL can serve as a fundamental educational device for AI practitioners and researchers, catalyzing experimentation across different communities.

5 DIAL-SYS Components

For interpreting DIAL it is assumed that users will have familiarity with elements of AI Systems, and as such there is an assumption that common acronyms are understood and need not be specified. As an example, “Bi-directional long-short term memory architecture” is written as BiLSTM without special introduction. Many components and symbols do require more definition: Table 2 contains an index to DIAL notation conventions.

Table 2: DIAL notational index

TABLE NUMBER	CONTENT AND ASSOCIATED CONVENTIONS
3	AI SHAPES, DOMAIN AND RANGE.
4	DIAL-SYS SYMBOLIC ELEMENTS.
5	DATA TYPES.
6	SOME EXAMPLE CLASSIFICATIONS.
7	DIAL-NN SYMBOLIC ELEMENTS.

5.1 Symbols

The iconic representation of DIAL, with its semiotic properties, provides the structuring framework for the diagrammatic representation. The DIAL-SYS *symbolset* provides a representation for high-level AI system components (as depicted in Table 3).

The primary DIAL-SYS *symbolset* can be extended with more specific symbolsets. Section 6 provides a description of DIAL-NN, a dialect aiming towards the standardisation of NN architectures.

Table 3: Shape inputs and outputs

SHAPE	DOMAIN	RANGE
POS TAGGING (POS)	S	S^{POS}
SYNTACTIC PARSING (SYN)	$S^{(POS)}$	PN STRUCTURE
NAMED ENTITY RECOGNITION (NER)	$S^{(Chunk)}$	S^{Names}
WORD SENSE DISAMBIGUATION (WSD)	$S^{POS,Chunk}, KB$	$Term^{WSD}$ OR WORDNET IDS
ENTITY LINKING (EL)	S^{NER}	ENTITY,(UPDATED KB)
SEMANTIC ROLE LABELING (SRL)	S^{Token}	S^{Sem}
SEMANTIC RELATION CLASSIFICATION (SRC)	$Term_1, Term_2$	$Term^{Sem}$
OPEN RELATION EXTRACTION (OIE)	S	$Pred(Arg)$
PREDICATE CREATION	$(T)(Data)$	$Term_{New} +$ DEF- INITIATION
STRUCTURED DATA QUERYING	$?q, KB$	TUPLES WITH LA- BELS
TEXT RETRIEVAL	HUMAN T	HUMAN- READABLE $T?$
NATURAL LANGUAGE GENERATION	$Pred(Arg)$	HUMAN- READABLE T
TEXT SIMPLIFICATION	T	T
TEXT SUMMARISATION	$T^{(Chunk)(NER)(Arg)(Name)(Sem)}$	T
CO-REFERENCE RESOLUTION (COREF)	S^{NER} OR $T^{(Token)}$	{CHAINS WITH IDS}
RHETORICAL STRUCTURE CLASSIFICATION (RST)	$(S);(S_1, S_2);(T)$	$S^{RS};T^{RS}$
ARGUMENTATION STRUCTURE CLASSIFICATION	$[s_1, s_2], T$	$T^{ArgStruct}$
ARGUMENT SCHEME CLASSIFICATION	$T^{ArgStruct}$	$C^{ArgScheme}$
POLARITY AND EMOTION ANALYSIS	$S^{(Pred(Arg))}$	SCORE
RHETORICAL FIGURES ANALYSIS	T	T^{RST}
STRING SIMILARITY	$string_1, string_2$	SCORE
SEMANTIC SIMILARITY	$terms^{(entity)}$	SCORE
SEMANTIC RELATEDNESS	$terms^{(entity)}$	SCORE
INDUCTIVE REASONING	$Pred(Arg)^F, (KB_R),$ $KB_{Constraints}$	$S^{(Pred(Arg))}$
DEDUCTIVE REASONING	$Pred(Arg)+KB_{F,R}$	$Pred(Arg)$
ABDUCTIVE REASONING	$Pred(Arg)^F, KB$	$Pred(Arg)$ SEQUENCE

Table 4: DIAL-SYS Symbolic Elements

FUNCTION	SYMBOL
DIRECT SUM	\oplus
CONCATENATION	\uplus
TENSOR PRODUCT	\otimes
SET	$\{elements\}$
DATA FLOW	\rightarrow
DATA FLOW (BOTH WAYS)	\leftrightarrow
KB QUERY	$\xrightarrow{?}$
DATA PERSISTENCE	\dashrightarrow
CONDITIONAL	\otimes
SYSTEM INTERFACE (E.G. SERVICE, API)	\dashv
COMPOSITION	$A \circ B$
JOIN	\bowtie
SIMILARITY & RELATEDNESS	$\angle\theta$
EMBEDDING PROJECTION (WITH IDENTIFIER)	$\bar{\Pi}_{id}$
WORD2VEC	$w2v$
SIMILARITY & RELATEDNESS (IF NOT COSINE, SPECIFY)	$\angle\theta$
REGRESSION	\boxed{R}
CLASSIFIER	\boxed{C}
CLASSIFICATION	C
RANKING OPERATOR	$R \uparrow$
TOP N ELEMENTS	$R \uparrow n$
ENCODER	\boxed{E}
DECODER	\boxed{D}
DEDUCTIVE REASONING	$E \models$
VERIFICATION (E.G. USER VALIDATION)	∇
FUNCTION	\boxed{f}
FUNCTION (CONTRACTION)	\xrightarrow{f}
DATASETS, DATA RESOURCES	$\boxed{\quad}$
GOLD STANDARD	\boxed{data}
KNOWLEDGE BASE OF FUNCTIONS	\boxed{Fn}
ZOOM IN	DOTTED LINE WITH BOX
ACCURACY	acc

5.2 Data Types & Notation Convention

DIAL symbols are typed with subscripts and superscripts, which describe recurring *classification tasks*, their associated *classes* or *data types*, tensor dimensionality, common data models (file formats). Table 2 indexes the tables. Features which form parts of the architecture are in *circles*, and full components are in *rectangles*. In this way, it is possible to see at a glance the level at which the system is abstracted.

Table 5: Data Types

CATEGORY	NOTATION	DESCRIPTION
TEXT	T	RAW FILE.
PASSAGE	p_T	FRAGMENT OF TEXT.
SENTENCE	s_T	SET OF WORDS FROM A TEXT, THAT IS COMPLETE IN ITSELF.
CHARACTER	Ch_T	A SINGLE PRINTABLE SYMBOL IN A TEXT.
TERM	t_T	A WORD OR SET OF WORDS IN A TEXT THAT EXPRESS A CONCEPT.
WORD	w_T	A SINGLE DISTINCT MEANINGFUL ELEMENT OF TEXT.
DIALOGUE	dt	A SET OF WORDS FORMING A DISTINT PART OF A DIALOGUE.
TERM		
SENSE	$\overline{w_T}$	A SPECIFICALLY DISAMBIGUATED WORD, INCLUDING ID.
CLUSTERED	\dot{w}_T	WORD EMBEDDED IN A VECTOR SPACE.
WORD		
IMAGE	im	RAW FILE.
SENTENCE	$\overline{s_T}$	AN IDENTIFIED SENTENCE, DISTINGUISHED FROM MINOR OR ADJACENT SENTENCES.
SENSE		
QUERY	$?q$	A SET OF WORDS AS A QUERY INPUT.
ANSWER	\hat{a}	A SET OF WORDS AS AN ANSWER OUTPUT.
FACTS	F	A PREDICATE WITH AN N-TUPLE OF WORD CONSTANTS, SATISFIED UNCONDITIONALLY.
RULES	R	A STATEMENT WHICH GIVES CONDITIONS UNDER WHICH TUPLES OF WORDS SATISFY A PREDICATE.
CLASSIFICATION	$P_c[a, b]$	THE PROBABILITY DISTRIBUTION OF A CLASSIFICATION c , WITH RANGE $[a, b]$.
OUTCOME		

6 DIAL-NN

As a second dialect, DIAL-NN (Neural Networks) was prioritised due to their prevalence in recent literature. As mentioned in the survey of Section 1.5, 77% of the papers surveyed were designing new Neural Networks or applying NN-based architectures.

Table 6: Example Classifications

CATEGORY	NOTATION
NER CLASSIFIED SENTENCE	S^{NER}
SRL CLASSIFIED SENTENCE	S^{SRL}
POS CLASSIFIED SENTENCE	S^{POS}
ARGUMENT SCHEME CLASSIFICATION	$C^{ArgScheme}$
ARGUMENTATION STRUCTURE	$T^{ArgStruct}$
CLASSIFIED AMBIGUOUS TERMS	$Term^{WSD}$
$Pred(Arg)$ LABELED WITH FACTS	$Pred(Arg)^F$

6.1 Symbols

DIAL-NN symbolic elements differ from DIAL-SYS symbolic elements due to distinguishing semiotic requirements. Table 7 shows the symbolic elements of this dialect.

Table 7: DIAL-NN Symbolic Elements

FUNCTION	SYMBOL
LOSS FUNCTION	Δ
ACTIVATION FUNCTION (LABEL WITH E.G. TANH)	
SOFTMAX	
ATTENTION	
RNN LAYER (EG LSTM)	
BiLSTM LAYER	
GRU LAYER	 GRU
CONVOLUTIONAL LAYER	
RECURSIVE NEURAL NETWORK	
SUPPORT VECTOR MACHINE	
GROUND TRUTH OF SENTIMENT CLASSIFICATION	g_c
HIDDEN LAYER (FORWARD DIRECTIONALITY)	\vec{h}
HIDDEN LAYER (BACKWARD DIRECTIONALITY)	$\leftarrow h$
...	...

7 DIAL by Example

7.1 Example 1: QA System

Figure 5 depicts an example DIAL-SYS diagram for a question answering system over unstructured text. The diagram depicts two cycles: (i) *KB construction*, which transforms textual data into a structured KB and (ii) *semantic parsing*, which transforms a natural language query into a formal structured query representation over the KB extracted in (i).

In the *KB construction* cycle, for each document d the pipeline iterates through all the sentences performing *open information extraction* (OIE) (as an external service), *named entity recognition* (NER) and *entity linking* (EL). This diagram does not describe in this context the OIE and NER steps, however it details the EL step. The EL starts by taking the complement to the NER chunks within the context of the sentence. In parallel, it collects the entity types associated to a structured KB (DBpedia). Both result sets are projected into a pre-built w2v space, where a cross-product between both sets is performed, cosine similarity measures are calculated and the top 1 ($argmax()$) is calculated. The output of the OIE component is joined to the output of the NER and EL, which is serialised into a RDF-NL file and indexed using an inverted index using the *tf-idf weighting scheme*. The RDF-NL file is composed with a pre-built word embedding, and this KB is exposed as a service.

In the *semantic parsing* cycle, a natural language question is sent to two parallel chains. In the upper part of the diagram, a POS tagger is applied, where the classifier type and the associated corpus are specified, as well as the accuracy in the original corpus and the accuracy in the current domain, this is then sent to a *lexical answer type* (LAT) detection component. In the lower part of the diagram, a C-Structure syntactic parsing is performed. Its output is added to the output of the POS tagger and sent to the SRL tagger. The LAT and the SRL are sent to a *Q-learning component*, which uses the KB of actions Op (queries over the structured KB - not detailed in this example) and the LAT + SRL, to learn a sequence of operations σ . The final output is given by the answer a .

Question Answering System

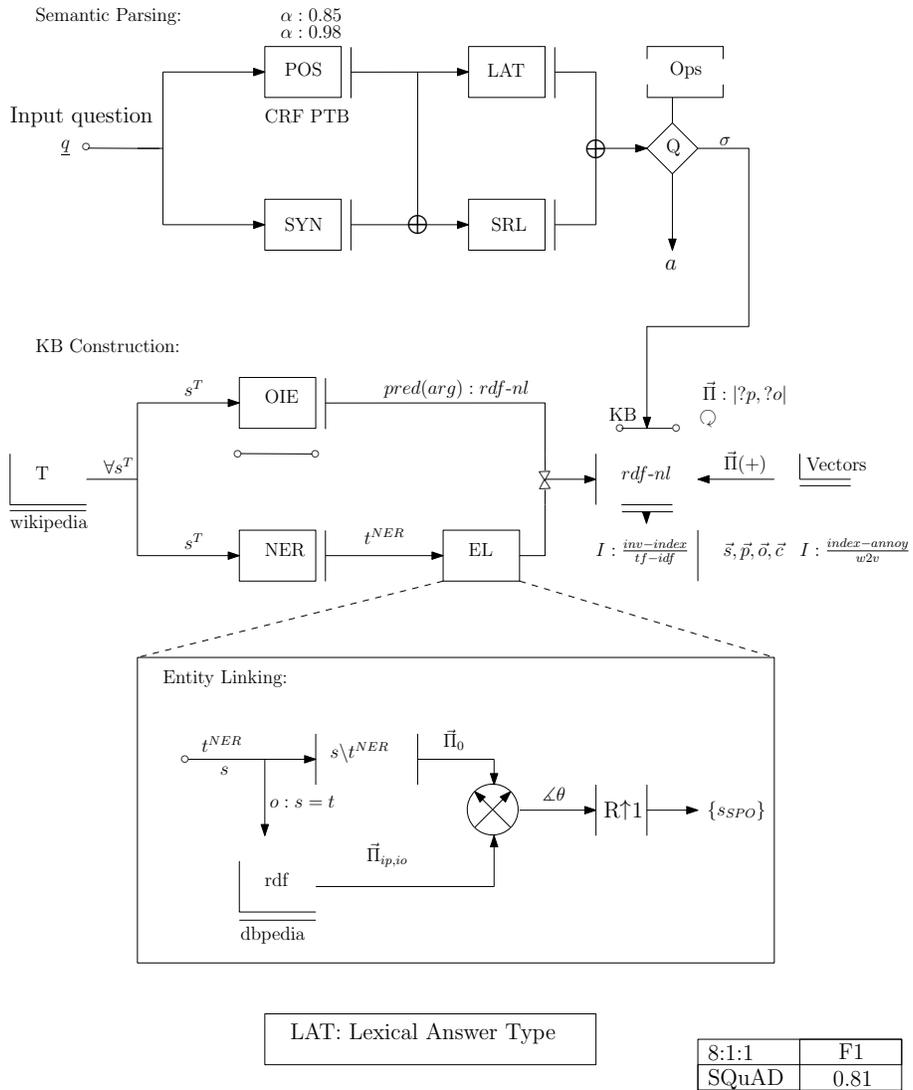


Figure 5: DIAL Diagram for a hypothetical QA system

7.2 Example 2: Lexicon-based Attention

The DIAL-NN dialect is illustrated by depicting the model described in Zou et al [17], which performs sentiment analysis over a document with attention at word-level and at the sentence level. Figure 6 depicts the organization of the architecture where the up most diagram describes the two Bi-LSTM layers with attention at the word-level and at the sentence level. Both models are jointly trained based on the creation of a word-level and sentence-level attention lexicon (lower left part). The hyper-parameters and the model accuracy are described in the two tables.

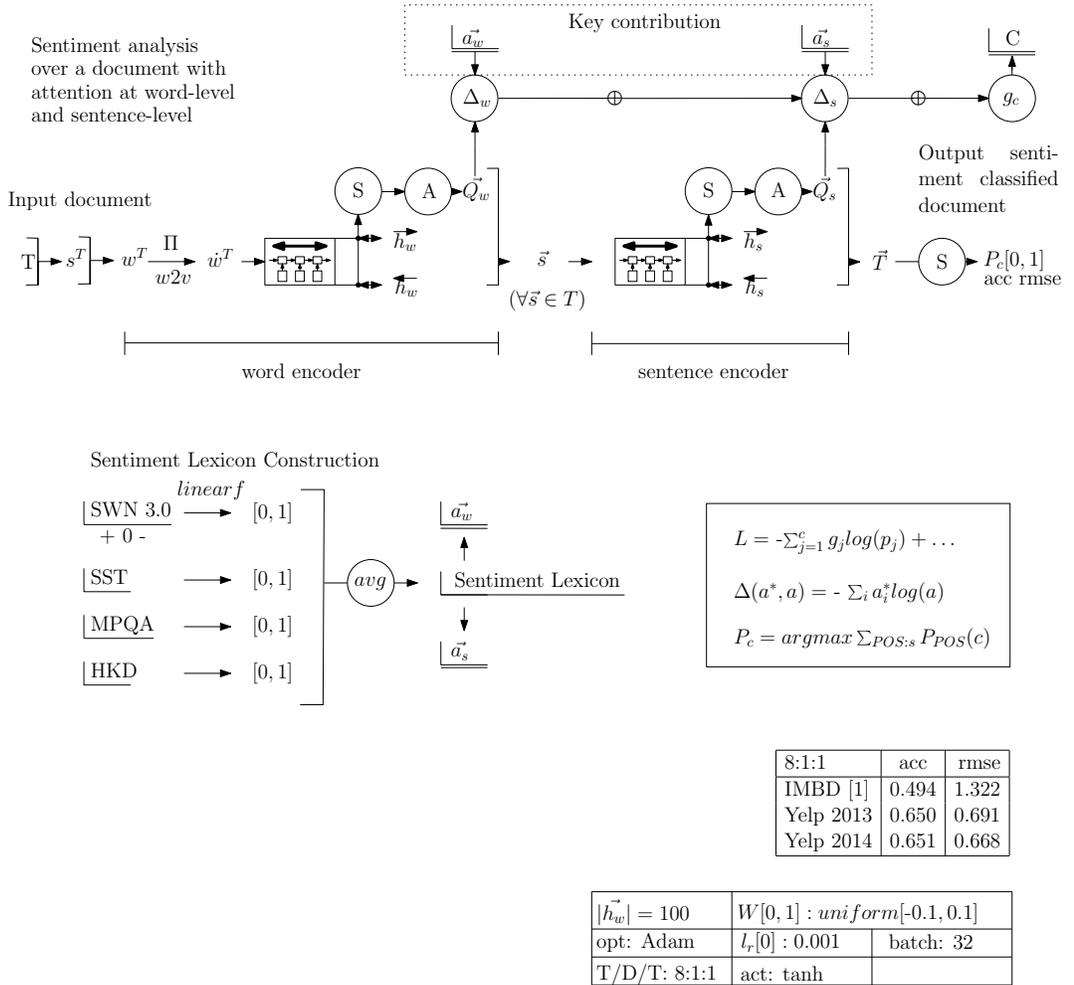
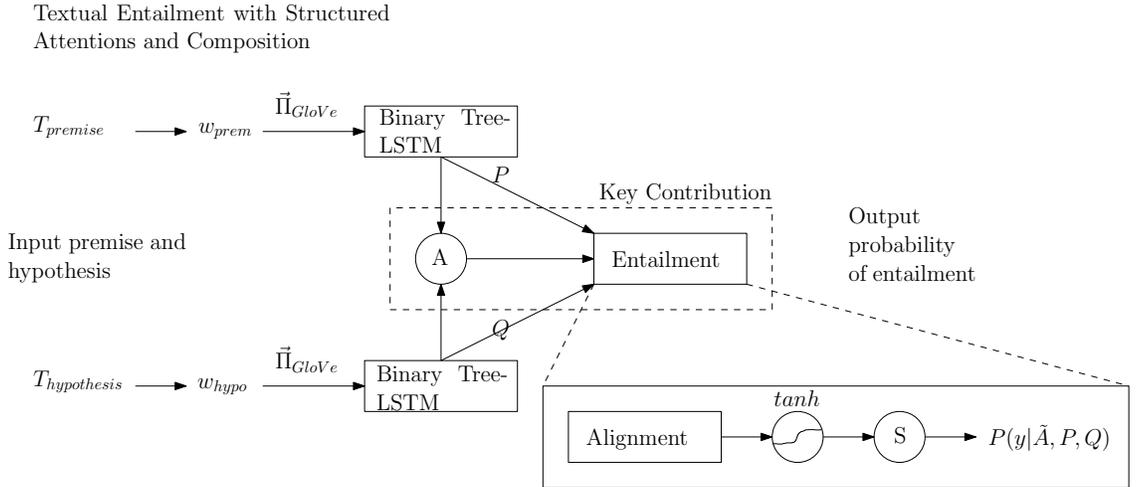


Figure 6: DIAL Diagram for Zou et al’s Lexicon-based supervision attention model

7.3 Example 3: Text Entailment

Figure 7 is a DIAL-SYS for Zhao et al [16]. Note that the details of the Neural Network are not included in this DIAL, though it could be extended to show this. Two Siamese projections (with shared weights) occur for premise and hypotheses, in order to create binary-tree LSTMs. A dual-attention model is applied to tree nodes between premise and hypothesis. The attention and outputs from binary trees for both premise and hypothesis are used in the entailment. At each hypothesis node i , e_i is calculated recursively given the meaning representation at this tree node h_i , the meaning representation of every node in the premise tree h_j , $j \in P$, and the entailment from i 's children, $e_{i,1}, e_{i,2}$. The entailment alignment is used in a tanh activation function for softmax, in order to output the probability of entailment between a premise and hypothesis. Note that this is a simplification due to forget gates etc.



$$\text{Entailment: } l \approx \frac{\exp(T_{2k,1}([h_i; h_j]))}{\sum_k \exp(T_{2k,1}([h_i; h_k]))}$$

$$\text{Alignment } \tilde{A}_{i,j} = Pr(A_{i,j} = 1 | P, Q)$$

$T_{2k,1}$ affine transformation R^{2k} to R

h_i calculated recursively from the meaning representations $h_{i,1}, h_{i,2}$ of its two children if node i is an internal node, otherwise h_i should be calculated based on the word $x \in R^d$ in the leaf. $h_i = f_{MR}(x_i, h_{i,1}, h_{i,2})$. Similar detail exists for relation e_i and LSTM transition

		SNLI
k	$ \theta _M$	F1
150	0.9m	87.2

Figure 7: DIAL Diagram for Textual Entailment with Structured Attentions and Composition [16]

8 Creating DIAL Diagrams

8.1 Practicalities

The DIAL examples contained in this report were made with IPE, a LaTeX compatible graphics editor (download source <http://ipe.otfried.org>). An accompanying DIAL resource library will be created in due course.

8.2 Semiotic Principles

By “semiotics”, we follow Morris in his definition of the field as encompassing semantics (what is represented), syntactics (relationships between things), and pragmatics (their interpretation within context) [10].

Spatial arrangement of symbols, arrows and natural language (Section 5) attempts to be reasonably standardised but not overly prescriptive. Additionally, they inherit existing common practices collected from the observation of the literature (see Survey section ??).

Top left title, bottom right data and hyperparameters are example of macroscopic guidelines. On a more detailed level, left-to-right linear is the default with exceptions where it makes sense (such as Siamese Networks, see Figure 7). Arbitrary alignment or misalignment is not encouraged, as they can confuse the representation. Making use of “intuitive” spatial patterns is also advisable, such as if the input is on the left, keep the input on the left in a “zoom in”. Any necessary functions, if necessary to detail, should be included in the mathematical key with notation consistent with the accompanying paper.

Introduce new symbols and relational concepts sparingly: We appreciate that DIAL is a work in progress and that some useful symbols will be missing. For example, it may be tempting to create a new symbol for linear scaling functions: Please avoid doing this, instead using existing natural language or mathematical notation. We are also aware that the practice of embedding text within an image is not ideal, and that there is a need for DIAL-NN to incorporate more low-level architectures (along the lines of Graves et al’s diagrams [6]). With interested community members we will be looking to adapt and extend DIAL in order to aid its usefulness. Please see Section 9 on how to become a contributor of DIAL.

9 Request for Comments (RfC)

Particularly at this early stage of development, feedback and interaction is sought from AI practitioners of all types. Feedback is especially helpful on topics that will aid completeness, consistency and usability of the existing scope of DIAL-SYS and DIAL-NN. As a longer-term commitment, those interested in being stewards for the standardisation effort should also get in touch with the authors. Further information, including contact details, can be found at aidiagrams.com.

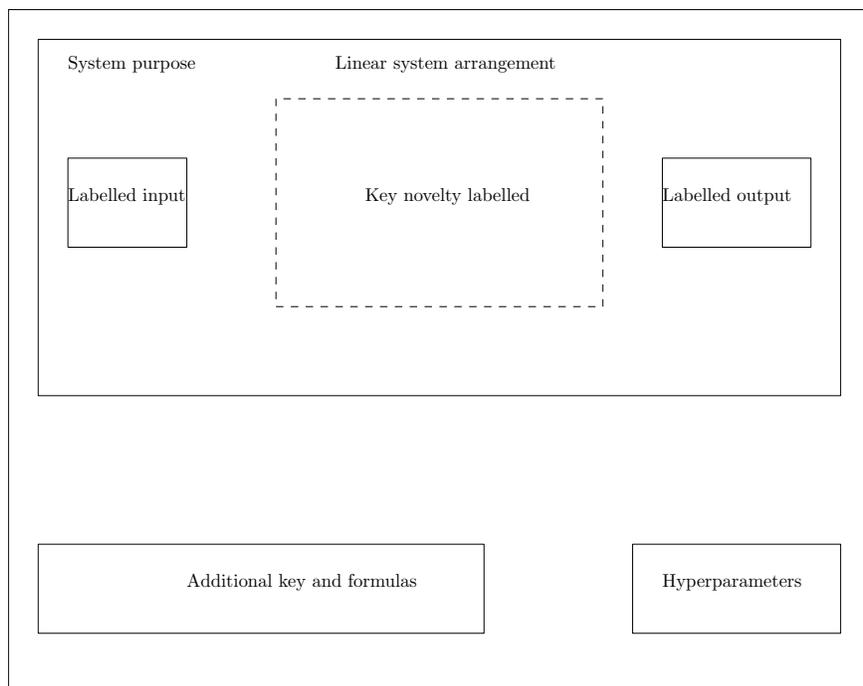


Figure 8: DIAL Schematic Outline

Acknowledgements

The authors would like to thank Viktor Schlegel, Chaohui Marshall and three anonymous reviewers for their useful comments.

References

- [1] ADEL, H., AND SCHÜTZE, H. Global Normalization of Convolutional Neural Networks for Joint Entity and Relation Classification. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing* (2017), Association for Computational Linguistics, p. 1723–1729.
- [2] AKHTAR, S., KUMAR, A., GHOSAL, D., EKBAL, A., AND BHATTACHARYYA, P. A Multilayer Perceptron based Ensemble Technique for Fine-grained Financial Sentiment Analysis. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing* (2017), pp. 540–546.
- [3] BOOCH, G., RUMBAUGH, J., JACOBSON, I., AND WESLEY, A. *Unified Modeling Language User Guide, The*. Addison Wesley, 1998.
- [4] EILENBERG, S., AND MACLANE, S. General theory of natural equivalences. *Transactions of the American Mathematical Society* 58 (1945), 231–231.
- [5] GLINZ, M. Problems and Deficiencies of UML as a Requirements Specification Language. In *Proceedings of the 10th International Workshop on Software Specification and Design* (2000), Institut für Informatik, Universität Zürich, IEEE Computer Society, pp. 11–22.
- [6] GRAVES, A., JAITLEY, N., AND MOHAMED, A.-R. Hybrid Speech Recognition with Deep Bidirectional LSTM. In *Automatic Speech Recognition and Understanding (ASRU), 2013 IEEE Workshop on* (2013), IEEE, pp. 273–278.
- [7] GUI, T., ZHANG, Q., HUANG, H., PENG, M., AND HUANG, X. Part-of-Speech Tagging for Twitter with Adversarial Neural Networks. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing* (2017), pp. 2411–2420.
- [8] HARRISON, N. The Darwin information typing architecture (DITA): applications for globalization. In *IPCC 2005. Proceedings. International Professional Communication Conference, 2005*. (2005), IEEE, pp. 115–121.
- [9] MOODY, D. What Makes a Good Diagram? Improving the Cognitive Effectiveness of Diagrams in IS Development. In *Advances in Information Systems Development*. Springer US, Boston, MA, 2007, pp. 481–492.

- [10] MORRIS, C. W. Foundations of the theory of signs. *The Journal of Symbolic Logic* 3, 04 (12 1938), 158.
- [11] OTERO, M. C., AND DOLADO, J. J. An empirical comparison of the dynamic modeling in OML and UML. *Journal of Systems and Software* 77, 2 (8 2005), 91–102.
- [12] SELBY, J. H., SCANDOLO, C. M., AND COECKE, B. Reconstructing quantum theory from diagrammatic postulates, 2018.
- [13] SETHI, A., SANKARAN, A., PANWAR, N., KHARE, S., AND MANI, S. DL-Paper2Code: Auto-generation of Code from Deep Learning Research Papers. In *Thirty-Second AAAI Conference on Artificial Intelligence* (2018), IBM Research, pp. 7339–7346.
- [14] SHAHIN, M., LIANG, P., AND BABAR, M. A. A systematic review of software architecture visualization techniques. *The Journal of Systems and Software* 94 (2014), 161–185.
- [15] SLOMAN, A. Why We Need Many Knowledge Representation Formalisms. In *Proceedings BCS Expert Systems* (1984), University of Sussex, Brighton, UK, Cambridge University Press, pp. 163–183.
- [16] ZHAO, K., HUANG, L., AND MA, M. Textual Entailment with Structured Attentions and Composition. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers* (2016), pp. 2248–2258.
- [17] ZOU, Y., GUI, T., ZHANG, Q., AND HUANG, X. A Lexicon-Based Supervised Attention Model for Neural Sentiment Analysis. In *Proceedings of the 27th International Conference on Computational Linguistics* (2018), pp. 868–877.