

A knowledge representation meta-model for rule-based modelling of signalling networks

Adrien Basso-Blandin¹, Walter Fontana², and Russ Harmer³

¹ LIP, ENS Lyon,
Lyon, France

`adrien.basso-blandin@ens-lyon.fr`

² Harvard Medical School,
Boston, USA

`walter@hms.harvard.edu`

³ CNRS & LIP, ENS Lyon,
Lyon, France

`russell.harmer@ens-lyon.fr`

Abstract

The study of cellular signalling pathways and their deregulation in disease states, such as cancer, is a large and extremely complex task. Indeed, these systems involve many parts and processes but are studied piecewise and their literatures and data are consequently fragmented, distributed and sometimes—at least apparently—inconsistent. This makes it extremely difficult to build significant explanatory models with the result that effects in these systems that are brought about by many interacting factors are poorly understood.

The rule-based approach to modelling has shown some promise for the representation of the highly combinatorial systems typically found in signalling where many of the proteins are composed of multiple binding domains, capable of simultaneous interactions, and/or peptide motifs controlled by post-translational modifications. However, the rule-based approach requires highly detailed information about the precise conditions for each and every interaction which is rarely available from any one single source. Rather, these conditions must be painstakingly inferred and curated, by hand, from information contained in many papers—each of which contains only part of the story.

In this paper, we introduce a graph-based meta-model, attuned to the representation of cellular signalling networks, which aims to ease this massive cognitive burden on the rule-based curation process. This meta-model is a generalization of that used by Kappa and BNGL which allows for the flexible representation of knowledge at various levels of granularity. In particular, it allows us to deal with information which has either too little, or too much, detail with respect to the strict rule-based meta-model. Our approach provides a basis for the gradual aggregation of fragmented biological knowledge extracted from the literature into an instance of the meta-model from which we can define an automated translation into executable Kappa programs.

1 Introduction

We propose a knowledge representation (KR) meta-model to enable the study of the *dynamics* of cellular signalling networks and, in particular, the consequences of mutations on dynamics. Our aim is therefore not to construct a Description Logic-based terminology (or any other ontology of this general kind) of static concepts to perform inference of the kind “ERK is an enzyme that phosphorylates so it is a kinase”; nor is it to build a representation *of* dynamics to support inference about the time-evolution of systems. Rather, we seek to represent each individual protein-protein interaction (PPI) that constitutes a signalling network as a formal *rule* that expresses the known, empirically necessary conditions for that PPI to occur.

These rules resemble those of Kappa or BNGL [11] but need not respect the stringent meta-model imposed by those formalisms wherein all bonds must occur between explicitly specified sites and all other relevant factors—such as protein conformation, post-translational modifications (PTMs) or, more generally, the presence or absence of key residues—are opaquely encoded into monolithic states attached to sites. In a second step, these rules can be automatically assembled into *bona fide* rule-based models that can be simulated, subjected to static analysis and whose causal structure can be examined in detail.

The need for such a two-stage approach to the rule-based modelling of signalling network-sarises because the pertinent information is dispersed across the literature in such a way that any given paper typically contains fragments, or *nuggets*, of partial mechanistic knowledge about multiple PPIs; and nuggets appear, for any given PPI, in many papers. Such a situation inevitably lends itself to a curation process focussed not directly on the PPIs themselves—as must any manual curation of rules—but rather on extracting nuggets, identifying which PPIs they refer to and incrementally aggregating them into more detailed nuggets.

A further novelty of our approach lies in our notion of agent which we take to represent not a single gene product but a *neighbourhood in sequence space*. This design decision is central to our approach and enables our KR to represent interactions that depend on certain key residues in such a way that the effects of both loss- and gain-of-function mutations can be automatically determined in the process of assembly to Kappa rules. This resolves a dangling question from earlier work on ‘meta Kappa’ [5, 10] which was only able to represent loss-, but not gain-, of-function mutations.

It is instructive to compare our approach with that embodied by production rule-based expert systems, such as MYCIN [1] and its descendants, which saw a gradual drift away from considering each rule as an independent element [as we do] towards a view where the KR should be “designed and maintained as a coherent whole” [2]. Such a viewpoint may be entirely appropriate in a domain where human experts can be reasonably expected to agree on most points; but, in a domain characterized by a large, dispersed and fragmented body of knowledge that no single human expert can hope to master, expert opinion can not reasonably be expected to converge towards a consensus.

As such, we advocate an approach where the KR does not seek to reproduce and augment human expert consensus but rather positions itself as a *tool for discovery* which, starting from purely objective nuggets of knowledge, enables and aids the human expert to investigate—and hopefully resolve—areas of apparent, or real, incoherence by comparing the dynamic consequences of various collections of independently-conceived rules. In particular, our system does *not* seek to figure out the ‘correct’ necessary conditions for a nugget and even less seeks to impose a pre-conceived structure on the Kappa model implied by the contents of the KR: there is no specification of what the system ‘should be’; indeed, the basic philosophy of our approach can be summarized as: taking as input partial empirical knowledge of a system; and producing as output the various consequences of this knowledge, *e.g.* the necessary conditions for PPIs, the causal structure, or pathways, that the system contains, *ℰc*. In a sense, from a computer science perspective, the workflow may seem backward: from partial knowledge of the behaviour of the system, we seek to *determine* its specification. Our approach is thus intrinsically oriented towards *systems*, not synthetic, biology which shares precisely this aim.

Overview of the paper. In section 2, we describe our graphical formalism, define our meta-model and formalize the notions of nuggets and their aggregation. In section 3, we illustrate these ideas with some simple examples. We conclude in section 4 with some remarks about our prototype implementation and directions for future work.

2 The graphical formalism

In this section, we first introduce a notion of (simple) graph with an additional forest structure on its nodes. This provides a more flexible and general starting point than in previous works [4, 6, 7] on ‘site graphs’ for rule-based modelling where the permitted kinds of nodes and edges were hard-wired. We then introduce a particular graph—our *meta-model*—which we use to *type* graphs so as to reintroduce the previously hard-wired constraints in a transparent fashion.

2.1 Site graphs

A *site graph* G is defined by

- a finite set of *nodes* \mathcal{N} with a specified subset of *states* \mathcal{S} ;
- a forest structure \mathcal{F} , *i.e.* a single-valued binary relation whose transitive closure is irreflexive, on \mathcal{N} such that all states are leaves and no states are roots;
- a simple directed edge structure \mathcal{E} , *i.e.* a binary relation, on \mathcal{N} ;
- a function assigning, to each state $s \in \mathcal{S}$, a set V_s of possible values.

The forest structure allows us to formalize the notion that a node may ‘belong to’ another, *e.g.* a node representing a binding site of some protein belongs to the node representing the protein in question. A state represents some (fixed or variable) *attribute* of a node, *e.g.* the identity of an amino acid at a certain sequence location or the presence or absence of a post-translational modification like phosphorylation; we sometimes refer to states of the latter kind as *flags*. The requirement that states be leaves (and not roots) enforces the idea that a state belongs to the node for which it acts as an attribute.

The ‘site graph’ terminology comes from a line of work on the definition of the rule-based modelling language Kappa in terms of graph rewriting [4, 6, 7] where various more constrained variants of this style of definition have been used. The principal novelty of the present definition is to avoid hard-wiring the various kinds of nodes that can exist—agents, sites, *ℰc.* —and their hierarchical structure. Instead, it proposes a homogeneous space of nodes with a forest structure that can capture arbitrary hierarchies.

2.2 Homomorphisms

A *homomorphism* $h : G \rightarrow G'$ of site graphs is a function $h : \mathcal{N} \rightarrow \mathcal{N}'$ such that states are preserved, *i.e.* $h(\mathcal{S}) \subseteq \mathcal{S}'$, the forest and its roots are preserved, edges are preserved and values are preserved, *i.e.* $V_s \subseteq V'_{h(s)}$ for all states s .

Site graphs and homomorphisms form a category **SGrph**; a homomorphism is a mono if, and only if, its underlying node function is injective. The category **SGrph** has all pull-backs, all push-outs and all pull-back complements over monos. As such, it possesses all the structure required to support general sesqui-push-out rewriting [3]. The sub-category of monos has all multi-sums [9].

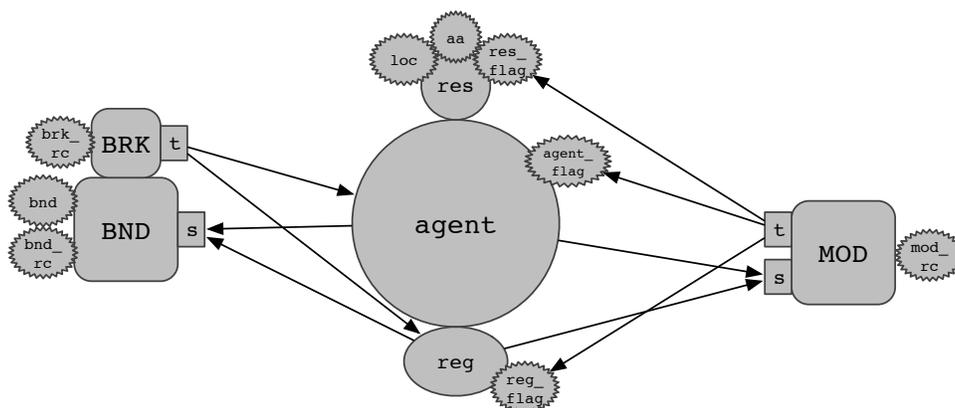
Given a fixed site graph G , the slice category over G , written **SGrph**/ G , can be usefully thought of as the category of graphs *typed by* G . Standard categorical reasoning establishes that the slice category construction preserves all the above categorical structure. In particular, the hierarchical structure of G and its edges constrain the objects of **SGrph**/ G so that we can reintroduce, in a flexible manner, the kinds of conditions and constraints that were hard-wired in previously proposed definitions of site graphs simply by selecting an appropriate base object G over which to take the slice category.

In section 2.3, we make such a choice for our knowledge representation by defining a specific graph \mathcal{M} which types, in the slice category sense, all possible models that we wish to consider, *i.e.* \mathcal{M} is our meta-model. In section 2.4, we introduce *nuggets* precisely as the class of site graphs that exist in $\mathbf{SGrph}/\mathcal{M}$. However, the general procedure of defining a desired class of site graphs by a choice of base object G could be applied in many other situations; it provides a unifying framework for discussing a broad class of related, but distinct, graphical formalisms.

2.3 The meta-model

We now introduce the concrete site graph \mathcal{M} that serves as our meta-model: it defines the various kinds of nodes that can exist, specifies which of them are states (and specifies the values they can take), defines the hierarchical structure on nodes and, finally, constrains the way that edges can be placed.

It is defined by the following graph where the forest structure is represented implicitly. The nodes labelled ‘agent’, ‘BND’ and ‘MOD’ are the roots; nodes with a ‘jagged’ outline are states.



We complete the definition by specifying the value sets of the states: `agent_flag`, `reg_flag`, `res_flag` and `bnd` are assigned the set $\{0, 1\}$ of Booleans; `loc`, which identifies the position of a residue in a sequence, is assigned the set \mathbf{Z}^+ of positive integers; `aa`, which identifies amino acids, is assigned the standard 20-element set of one-letter amino acid codes, *i.e.* all letters of the alphabet except B, J, O, U, X and Z; and `bnd_rc`, `brk_rc` and `mod_rc`, which specify the rate constants of actions, are assigned the set \mathbf{R}^+ of positive reals.

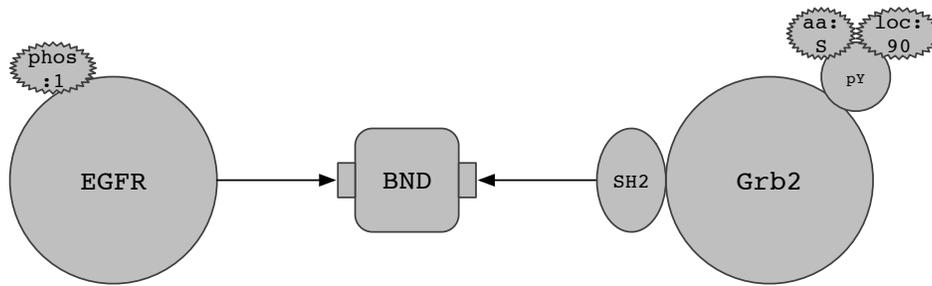
Note that the names of nodes, such as ‘agent’ and ‘BND’, are not part of the formalism; they are just a convenient labelling for the purposes of discussion.

A typing $h : G \rightarrow \mathcal{M}$ formally designates each of the nodes of G as being either an *agent*, a *region*, a *residue*; or as one of the fixed attributes (`aa`, `loc`, *etc.*); or as a *flag*; or as a *binding* action (BND), *unbinding* action (BRK), *modification* action (MOD) or as a *source* for a BND or MOD or as a *target* for a BRK or MOD. The mapping of G into \mathcal{M} further implies that the forest structure of G must respect the restrictions imposed by \mathcal{M} , *i.e.* that regions and residues belong to agents; that attributes and flags must be leaves; that sources and targets belong to their respective actions; and that unbinding actions belong to their corresponding binding actions. Finally, the edge structure of G must also respect the restrictions imposed by \mathcal{M} ; this means that only agents and regions can engage in binding actions; and that only flags can be targeted by modification actions.

This meta-model provides the foundation for a rigorous ontology for the kinds of information that are pertinent to rule-based descriptions of signalling networks, *i.e.* proteins as agents, domains and other binding sites as regions, key amino acid locations as (key) residues, *ℳc*. However, the present framework remains purely formal and does not have any means to enforce correct semantic usage of this ontology; as such, we plan to augment our framework with a system of appropriate annotations in order to be able to carry out semantic checking and reasoning. We will return to this point later.

2.4 Nuggets

Let us now motivate the particular choice \mathcal{M} of meta-model made in the previous section by considering a typical ‘nugget’ of knowledge in molecular biology: “EGFR binds the SH2 domain of Grb2 provided that EGFR is phosphorylated and residue 90 of Grb2 is a serine”. This would naturally be represented as the following site graph



where the agent ‘EGFR’ has a state ‘phos’ (with value 1, meaning true) and the agent ‘Grb2’ has a region ‘SH2’ and a residue (with unimportant name) located at position 90 of the sequence and which is required to be a serine (the value S). Note that no region has been specified on EGFR; the meta-model explicitly allows for this by the fact that a BND action can have either an agent or a region as its source.

Formally speaking, a *nugget* is a connected graph G typed by \mathcal{M} in such a way that each node has at most one copy of each of its possible attributes, where all attributes and flags have a uniquely specified value and exactly one action node does not have a specified value for its `bnd` flag. This unique action is the *principal* action of the nugget; any other (necessarily BND) actions represent the required presence or absence of their corresponding bonds in order for the principal action to be possible. We further ask that any BND has exactly two sources, any BRK has exactly two targets and any MOD has at most one source and one target (and at least one of the two). Other than asking for G to be connected, all of these constraints are motivated by domain-specific considerations and, in due course, will be handled via semantic checking.

Our representation language is thus a generalization of that typically used in rule-based modelling. The principal differences are that (i) the *action* of a rule is represented explicitly as a node in the graph; (ii) binding actions can act directly on agents rather than necessarily via sites/regions; and (iii) static attributes and dynamic flags are represented as values associated to persistent nodes rather than as mutually exclusive sets of nodes. We draw a distinction between flags and attributes in order to make an explicit difference between things that can be modified ‘in the system’, *e.g.* the phosphorylation status of some residue, and things that can only be modified ‘out of the system’, *e.g.* the identity of an amino acid that can only be modified by an external mutation event.

The latter points (ii) and (iii) are important in order to be able to represent biological knowledge as faithfully as possible: knowledge is often stated in a piecemeal and incomplete fashion but this should not prevent us from being able to formalize it. For example, when the site at one end of a bond is unknown, this can now be represented *as is* with no need—as there would have been in standard rule-based modelling—to create a ‘fictitious’ site. A second example could be the use of attributes to provide a transparent representation of detailed structural information, *e.g.* about key residues of a protein or cases where a binding interaction depends on multiple PTMs that can otherwise only be opaquely encoded. We will return to point (i) in section 2.6 where we will introduce the notion of nugget *aggregation* which depends critically on the explicit representation of actions as nodes.

2.5 Models

A site graph $m : M \rightarrow \mathcal{M}$ is a *pre-model* of a collection C of nuggets $n_i : N_i \rightarrow \mathcal{M}$ iff, for all i , the arrow n_i factors through m . In words, M is a graph, itself typed by the meta-model, which types all of the nuggets; indeed, M can be thought of as a summary statement of the collection C of nuggets. We refer to the pair (C, M) as a *model*. Note that a given collection of nuggets may be assigned many different pre-models. The import of any particular choice is that it identifies which nodes in one nugget correspond to those in another: in the above example, we have a node labelled as ‘EGFR’; but that label does not exist in the formalism so, if we have a second nugget which also speaks of the same agent ‘EGFR’, we need a way to say that these two nodes are the same. The pre-model gives us precisely this possibility: the two agents are mapped to the same node of M , *i.e.* nodes of M provide labels/names for the nodes of nuggets. This means that two different pre-models, M_1 and M_2 , for the same collection C of nuggets can have completely different meanings; in particular, C and \mathcal{M} provide two (uninteresting) extremes with interesting cases lying in between. In general, the two components of a model evolve together as we add more and more information; we discuss this briefly in section 2.6.

Let us note here that the necessity of a pre-model partially arises in order to enforce minimal semantic coherence in our formal framework. If we had semantic annotations that uniquely identify agents, we could potentially use them—instead of a pre-model—to solve the above cross-nugget identification problem. However, we have chosen to take a different approach so as to provide a more flexible notion of agent: in a general site graph—and, in particular, in a pre-model—the **aa** attribute of a residue may be assigned a *set* of one-letter codes in order to express the fact that an agent represents, in general, a *neighbourhood* in sequence space rather than a unique sequence. This flexibility affords us the possibility of organizing knowledge about minor variants of a protein using a single agent; this (rather prosaically) pre-empts the need to define, and name, lots of tedious variants but, more to the point, matches everyday practice in biology where, for example, (wild-type) ‘Ras_{WT}’ and (mutant) ‘Ras_{G12V}’ are both thought of as being ‘Ras’—they just differ in one or two small, although possibly very significant, ways.

As such, the notion of model is our first step towards a full semantic layer for our knowledge representation scheme: all semantic annotations will be made at the level of pre-models, not individual nuggets, so as to minimize the amount of needed annotation and, more importantly, to ease the maintenance of semantic coherence across the entire current collection of nuggets. It should be noted that this approach to grounding differs from more traditional approaches, such as that used by BioPAX [8], which insist upon each formal entity corresponding to a unique physical entity. Indeed, our approach is, by design, particularly attuned to the needs of representing signalling networks and, as such, is less constrained than BioPAX which, as a framework of far broader applicability, has to bear a far stronger semantic burden.

2.6 Aggregation

Suppose that, at some point in time, we have a model (C, M) and that we now obtain a new nugget N' . We always have the possibility simply to add N' , yielding a new collection C' of nuggets; this might necessitate updating the underlying pre-model M to M' in the event that N' contains entirely novel nodes or edges.

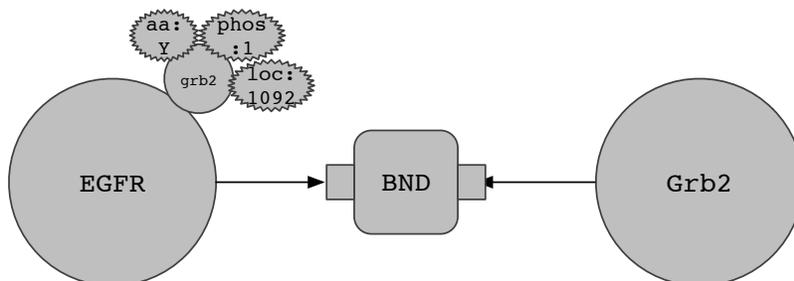
If there was already a nugget in C that has the *same* action as that of N , this would result in (C', M') having two distinct actions involving the same agents, a situation that may or may not be desirable: sometimes, for example, two proteins can indeed bind each other in two distinct ways; however, it could also be the case that N' has actually brought some new information about a single binding interaction that we would rather use to *update* the pre-existing nugget N . Such an update amounts to the assumption that the nuggets N and N' represent the same interaction *mechanism*; but this does not necessarily mean that the two nuggets refer to exactly the same agents since mechanisms can be shared across families of proteins. We will return to this point in section 3 after briefly describing the formal process of *aggregation*.

If we wish to *update* the nugget N with the information contained in N' , we need to specify two things: the new information brought by N' and any *deprecated* information in N that should now be removed. The former is specified by the choice of a co-span of monos $h_+ : N \rightarrow N_+ \leftarrow N' : h'_+$ from the multi-sum of N and N' ; while the latter is specified by a mono $h_- : N_- \rightarrow N$. In most cases, there is a canonical choice of co-span given by the intuitive unification of N and N' but, in cases where there are non-trivial automorphisms of N and N' , the unification process may be non-deterministic and a choice becomes necessary (or we allow multiple conflicting versions).

The pull-back complement of h_- and h_+ defines a graph N_{\pm} containing precisely the new information from N' with all deprecated information from N removed. Formally, this is exactly a step of graph rewriting taking N to N_{\pm} . In the event that nothing is to be removed from N , *i.e.* h_- is the identity on N , this rewriting step degenerates to being simply the refinement of N to N_+ as specified by h_+ . This step of rewriting is also propagated to the pre-model M , resulting in a new model (C', M') where, unlike in the case of adding N' , C' clearly has the same number of nuggets as C .

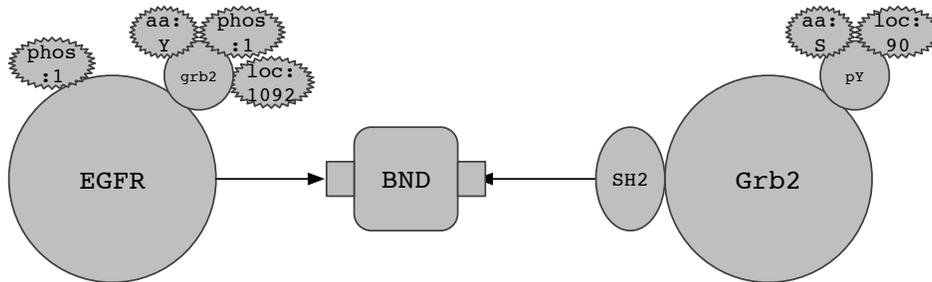
3 Examples of aggregation

Nugget update Consider updating the example nugget N of section 2.4 with the information contained in N' :

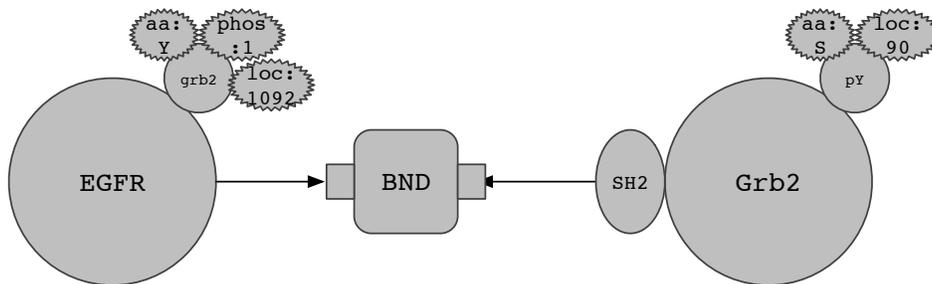


i.e. “EGFR binds Grb2 provided EGFR is phosphorylated on Y1092”.

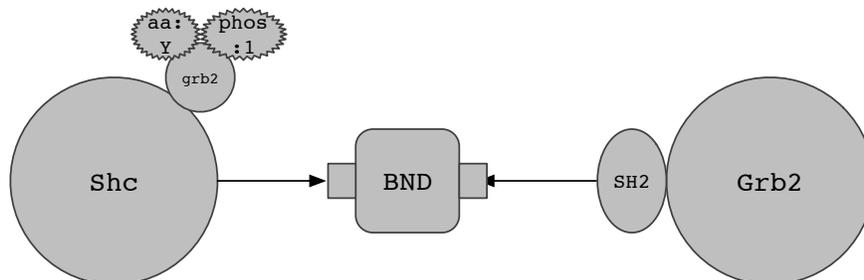
If we choose not to deprecate anything from the original nugget, we obtain:



If instead we were to specify that the **phos** flag from *N* is to be removed, we would obtain:

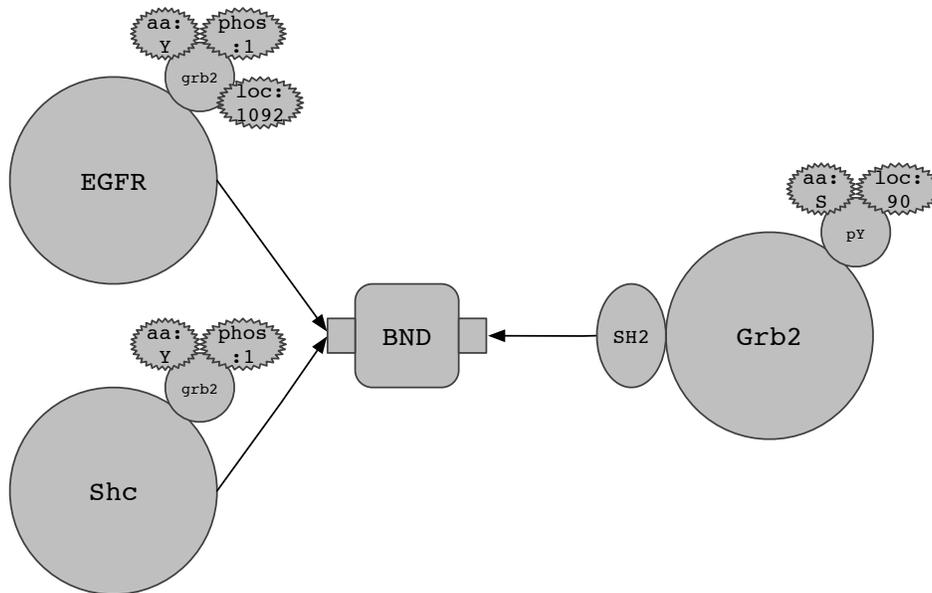


Nugget aggregation Unification can also be partial: if we further update with



i.e. “tyrosine-phosphorylated Shc binds the SH2 domain of Grb2”

we obtain:



Note how the contextual conditions on **Grb2** that occur in the original nugget are propagated, by the very process of aggregation, to its newly added interaction with **Shc**. This is a typical example of the use of our framework as a tool for discovery at the level of necessary conditions for PPIs which renders completely transparent the ‘by similarity’ style of reasoning which is ubiquitous in molecular biology.

Note that **EGFR** and **Shc** both target the same source of the (unique) **BND** node: this nugget has a disjunctive interpretation, giving rise to two distinct Kappa rules, and should be read (modulo the contextual conditions) as “EGFR or Shc binds the SH2 domain of Grb2”. The importance of this is that, in the translation to Kappa, agents are assigned one *site* for each **BND** action in which they participate: before aggregation, **Grb2** would have been given two sites, one to bind **EGFR** and the other for **Shc**, with no conflict between the two generated Kappa rules; after aggregation, it would receive only one site, giving rise to an intrinsic conflict between the two generated rules.

4 Conclusion

We have presented a meta-model for aggregation of the kind of knowledge required to build rule-based models of cellular signalling networks. The framework provides a generalization of the usual strict rule-based meta-model and, in particular, represents the actions of rules explicitly as nodes. This enables the key notion of aggregation which serves as our source of biologically-plausible inference. The presentation is entirely mathematical, being framed in terms of a graph rewriting formalism, but can be considered as a specification for an actual system. We are currently building a prototype implementation of the framework, including the automatic translation of a collection of nuggets into *bona fide* Kappa, and will report on this in more detail in a full-length version of this paper.

Let us nonetheless discuss one important aspect of the translation: each BND node generates its own sites; so an agent with two incident BND actions will be translated into a Kappa agent (let's call it A) with two sites. This is a consequence of our general philosophy that, in the absence of information to the contrary, we should draw the most general conclusions possible: were we to generate but a single site for that Kappa agent A , we would be unwarrantedly hard-wiring the constraint that its binding partners must compete in order to bind. On the other hand, in the event that we subsequently learn that both BND actions depend on some common residue of A , this would imply an intrinsic conflict between the two actions and would result in the automatic translation generating a single binding site, thus enforcing competitive binding to A . (This is a more general manifestation of the above example of nugget aggregation where the fusing of two actions gave rise to an analogous conflict.) In this way, we achieve a pragmatically (and cognitively) convenient separation of concerns whereby knowledge is integrated and aggregated as we learn it; while its eventual consequences for a future Kappa model are determined, at model generation time, by the automatic translation procedure.

After the automatic translation into Kappa, each concrete agent that is generated must have a unique value for each of its `aa` attributes; these manifest as *unmodifiable* states in the resulting Kappa. As such, a nugget that tests such an `aa` attribute will be translated into a Kappa rule (or rules) that will only ever apply to those concrete agents that have the appropriate state. As such, a rule that tests for a wild-type value of such an attribute would only apply to a concrete agent that is not mutated at that residue; conversely, a rule that tests for a non-wild-type value would not apply to a concrete agent unless it had undergone an appropriate mutation. This therefore enables a transparent account of both loss- and gain-of-function mutations that side-steps the various difficulties (concerning gain-of-function mutations) that were encountered by the original 'meta Kappa' project [5, 10]

Finally, the framework presented here remains entirely formal and, although we obviously have in mind an interpretation in terms of signalling networks, it does not actually embody any domain-specific knowledge. In particular, although we have used suggestive names for nodes in our examples, these have no actual significance and could be arbitrarily renamed without affecting the content of the knowledge representation. We plan to address issue this by introducing semantic annotations for nodes that would allow us to express domain-specific properties such as "this region is an SH2 domain" or "this residue must be serine or threonine".

This *grounding* process opens up the possibility of performing a second level of semantic checking on nuggets that have already been verified to be syntactically well-formed. Such checks could be purely routine, *e.g.* "binding actions have exactly two participants", "only a kinase can phosphorylate a protein" or "a serine/threonine kinase cannot phosphorylate a tyrosine residue". However, we principally envisage semantic annotations as a means to *automate* certain default aggregation decisions, *e.g.* "an SH2 domain can bind only one phospho-tyrosine ligand at a time", with a particular focus on the numerous domain-domain and domain-ligand PPIs that occur in signalling networks which, by their very nature, embody highly generic binding mechanisms constrained by relatively simple—but tedious and error-prone to write by-hand—conflict conditions.

Acknowledgement. This work was sponsored by the Defense Advanced Research Projects Agency (DARPA) and the U. S. Army Research Office under grant number W911NF-14-1-0367. The views, opinions, and/or findings contained in this report are those of the authors and should not be interpreted as representing the official views or policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the Department of Defense.

References

- [1] Bruce G Buchanan and Edward H Shortliffe. *Rule-based Expert Systems: the MYCIN Experiments of the Stanford Heuristic Programming Project*. Addison-Wesley, 1984.
- [2] William J Clancey. Notes on “Epistemology of a rule-based expert system”. *Artificial Intelligence*, 59(1):197–204, 1993.
- [3] Andrea Corradini, Tobias Heindel, Frank Hermann, and Barbara König. Sesqui-pushout rewriting. In *Graph Transformations*, pages 30–45. Springer, 2006.
- [4] Vincent Danos, Jérôme Feret, Walter Fontana, Russ Harmer, Jonathan Hayman, Jean Krivine, Chris Thompson-Walsh, and Glynn Winskel. Graphs, rewriting and pathway reconstruction for rule-based models. In *FSTTCS 2012-IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science*, volume 18, pages 276–288, 2012.
- [5] Vincent Danos, Jérôme Feret, Walter Fontana, Russ Harmer, and Jean Krivine. Rule-based modelling and model perturbation. In *Transactions on Computational Systems Biology XI*, pages 116–137. Springer, 2009.
- [6] Vincent Danos, Russ Harmer, and Ricardo Honorato-Zimmer. Thermodynamic graph-rewriting. In *CONCUR 2013-Concurrency Theory*, pages 380–394. Springer, 2013.
- [7] Vincent Danos, Russ Harmer, and Glynn Winskel. Constraining rule-based dynamics with types. *Mathematical Structures in Computer Science*, 23(02):272–289, 2013.
- [8] Emek Demir, Michael P Cary, Suzanne Paley, Ken Fukuda, Christian Lemer, Imre Vastrik, Guanming Wu, Peter D’Eustachio, Carl Schaefer, Joanne Luciano, et al. The biopax community standard for pathway data sharing. *Nature biotechnology*, 28(9):935–942, 2010.
- [9] Y. Diers. Familles universelles de morphismes. *Tech. report, Université des Sciences et Techniques de Lille I*, 1978.
- [10] Russ Harmer. Rule-based modelling and tunable resolution. *arXiv preprint arXiv:0911.2508*, 2009.
- [11] William S Hlavacek, James R Faeder, Michael L Blinov, Richard G Posner, Michael Hucka, and Walter Fontana. Rules for modeling signal-transduction systems. *Science’s STKE*, 2006(344):re6, 2006.