

Interpretable exact linear reductions via positivity^{*}

Gleb Pogudin¹ and Xingjian Zhang²

¹ LIX, CNRS, École Polytechnique, Institut Polytechnique de Paris, Palaiseau, France gleb.pogudin@polytechnique.edu

² École Polytechnique, Institut Polytechnique de Paris, Palaiseau, France xingjian.zhang@polytechnique.edu

Abstract. Kinetic models of biochemical systems used in the modern literature often contain hundreds or even thousands of variables. While these models are convenient for detailed simulations, their size is often an obstacle to deriving mechanistic insights. One way to address this issue is to perform an exact model reduction by finding a self-consistent lower-dimensional projection of the corresponding dynamical system. Recently, a new algorithm CLUE [16] has been designed and implemented, which allows one to construct an exact linear reduction of the smallest possible dimension such that the fixed variables of interest are preserved. It turned out that allowing arbitrary linear combinations (as opposed to zero-one combinations used in the prior approaches) may yield a much smaller reduction. However, there was a drawback: some of the new variables did not have clear physical meaning, thus making the reduced model harder to interpret.

We design and implement an algorithm that, given an exact linear reduction, re-parametrizes it by performing an invertible transformation of the new coordinates to improve the interpretability of the new variables. We apply our algorithm to three case studies and show that “uninterpretable” variables disappear entirely in all the case studies.

The implementation of the algorithm and the files for the case studies are available at <https://github.com/xjzhang/LumpingPostiviser>.

Keywords: exact reduction (lumping) · ODE model · interpretability.

1 Introduction

Dynamical models described by systems of polynomial ordinary differential equations (PODEs) are frequently used in systems biology and life sciences in general. One of the major classes of such models is the dynamical models of chemical reaction networks (CRN) under the mass-action kinetics in which each indeterminate corresponds to the concentration of one of the chemical species. Models

^{*} Supported by the Paris Ile-de-France region. GP was partially supported by NSF grants DMS-1853482, DMS-1760448, DMS-1853650, CCF-1564132, and CCF-1563942.

appearing in the literature often consist of hundreds or thousands of variables. While the models of this size can incorporate a substantial amount of information about the phenomena of interest, it is often hard to use them to derive mechanistic insights.

One way to address these challenges is to use *model reduction* algorithms that replace a model with a simpler one while preserving, at least approximately, some of the features of the original model. A wide range of methods has been developed for approximate model reduction, including methods based on singular value decomposition [1] and time-scale separation [15].

A complementary approach is to perform *exact model reduction*, that is, lower the dimension of the model without introducing approximation errors. For example, exact linear lumping aims at writing a self-consistent system of differential equations for a set of *macro-variables* in which each macro-variable is a linear combination of the original variables. For important classes of biochemical models, specialized lumping criteria have been developed (see, e.g., [3,6,9]), allowing the construction of macro-variables as sums of some of the original variables (that is, allowing only coefficients zero and one in the linear combinations). A general lumping algorithm has been proposed in [5,4] which is applicable to any system of PODEs (not necessarily arising from a CRN). This algorithm partitions the original variables so that the macro-variables can be the sums of the variables within the blocks in the partition. Note that the macro-variables are zero-one linear combinations of the original variables in all these cases.

In [16], an algorithm has been designed (and the corresponding software called *CLUE* presented) that, for a given set of linear forms in the state variables (the *observables*), constructs a linear lumping of the smallest possible dimension such that the observables can be written as combinations of the macro-variables (i.e., the observables are preserved). Unlike the earlier approaches, the macro-variables produced by CLUE may involve any coefficients, and this allowed to produce reductions of lower dimensions than it was possible before, see [16, Table 1]. However, there was a price to pay for this flexibility: the authors state that some of the produced macro-variables “escape physical intelligibility” (see [16, Section 4.2]). Indeed, the resulting reduction of the smallest dimension is uniquely defined up to a linear change of the coordinates, so the coordinates in the reduced state space chosen by CLUE could be not optimal in the sense of interpretability.

In this paper, we propose a post-processing step that takes an exact linear lumping (not necessarily produced by CLUE) and attempts to improve its interpretability by performing a change of variables. It has been observed in [16] that one of the sources of difficulties for interpretation is the negative coefficients in the macro-variables. We design and implement an algorithm that finds (if possible) a linear change of variables in the reduced model so that

1. the coefficients of the representations of the new macro-variables in terms of the original state variables are nonnegative
2. and the total number of nonzero coefficients in these representations is as small as possible.

Note that interpretability is not a formal mathematical property, and the conditions above is one possible formalization of the notion of a “more interpretable reduction”. We do not claim that it is universal (e.g., a difference of two state variables may represent a potential), but we claim that it is useful. To support this claim, we demonstrate the efficiency of our approach on three case studies from the literature. Two of these cases are exactly the case studies from [16] in which issues with interpretability occur. We show that our method provides interpretable re-parametrizations of the optimal lumpings computed by CLUE in all three case studies. Our algorithm uses tools from convex discrete geometry and matroid theory.

2 Methods

2.1 Preliminaries on lumping

Definition 1 (Lumping). *Consider a system of ODEs of the form*

$$\mathbf{x}' = \mathbf{f}(\mathbf{x}), \quad (1)$$

where $\mathbf{x} = (x_1, \dots, x_n)^T$, $\mathbf{f} = (f_1, \dots, f_n)^T$, and $f_1, \dots, f_n \in \mathbb{R}[\mathbf{x}]$. A linear transformation $\mathbf{y} = L\mathbf{x}$ with $\mathbf{y} = (y_1, \dots, y_m)^T$, $L \in \mathbb{R}^{m \times n}$, and $\text{rank } L = m$ is called a lumping of (1) if there exist polynomials $g_1, \dots, g_m \in \mathbb{R}[\mathbf{y}]$ such that

$$\mathbf{y}' = \mathbf{g}(\mathbf{y}), \quad \text{where} \quad \mathbf{g} = (g_1, \dots, g_m)^T$$

for every solution $\mathbf{x}$ of (1). We say that m is the dimension of the lumping. The variables $\mathbf{y}$ in the reduced system are called macro-variables. We will call a macro-variable nontrivial if it is not proportional to one of the original variables.

Remark 1. An ODE system may have many lumpings, some of them may be less useful than others. For example, if $m = n$, then the lumping is just an invertible change of variables, so no reduction happens. Another special case is when the rows of L contain the coefficients of linear first integrals of the system. In this case, the reduced ODE will be of the form $\mathbf{y}' = 0$.


Constrained linear lumping introduced in Definition 2 requires to preserve the dynamics of the variables of interest, and this is one of the ways to say that reduction is not “too coarse”.

The following example is a substantially simplified version of the case study from Section 3.1 (see also [12]).

Example 1. We will consider a chemical reaction network consisting of

- A chemical species X .
- Species A_{UU} , A_{UX} , A_{XU} , and A_{XX} . Each of them is one of the states of a molecule A with two identical binding sites, which can be either unbound (U in the subscript) or bound (X in the subscript) to X .

For simplicity, we will assume that all the reaction rates are equal to one. The dynamics of the network is defined by the following reactions (* denotes any of X and U):



Under the laws of the mass-action kinetics, the reactions (2) yield the following ODE system (where $[S]$ denotes the concentration of the species S):

$$\begin{cases} [X]' = [A_{XU}] + [A_{UX}] + 2[A_{XX}] - [X]([A_{XU}] + [A_{UX}] + 2[A_{UU}]), \\ [A_{UU}]' = [A_{XU}] + [A_{UX}] - 2[X][A_{UU}], \\ [A_{XU}]' = [A_{XX}] + [X][A_{UU}] - [X][A_{XU}] - [A_{XU}], \\ [A_{UX}]' = [A_{XX}] + [X][A_{UU}] - [X][A_{UX}] - [A_{UX}], \\ [A_{XX}]' = [X][A_{XU}] + [X][A_{UX}] - 2[A_{XX}]. \end{cases} \quad (3)$$

We will show that the following matrix L and the macro-variables y_1, y_2, y_3

$$L = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 2 \\ 0 & 2 & 1 & 1 & 0 \end{pmatrix} \implies \begin{cases} y_1 = [X], \\ y_2 = [A_{XU}] + [A_{UX}] + 2[A_{XX}], \\ y_3 = 2[A_{UU}] + [A_{XU}] + [A_{UX}]. \end{cases} \quad (4)$$

yield a lumping of the system (2). Indeed, a direct calculation shows that

$$\begin{cases} y_1' = [X]' = [A_{XU}] + [A_{UX}] + 2[A_{XX}] - [X]([A_{XU}] + [A_{UX}] + 2[A_{UU}]) = y_2 - y_1 y_3, \\ y_2' = [A_{XU}]' + [A_{UX}]' + 2[A_{XX}]' = -y_2 + y_1 y_3, \\ y_3' = 2[A_{UU}]' + [A_{XU}]' + [A_{UX}]' = y_3 - y_1 y_2. \end{cases} \quad (5)$$

Since each reaction involves only one binding site, this lumping can be interpreted as follows: y_2 is the total “concentration” of the bound sites, and y_3 is the total “concentration” of the unbound sites (see also Section 3.1).

The lumping matrix L in the example above turns out to exactly preserve the concentration $[X]$. In general, one may fix a vector $\mathbf{x}_{\text{obs}}$ of combinations of the original variables that are to be recovered in the reduced system.

Definition 2 (Constrained linear lumping). *Let $\mathbf{x}_{\text{obs}}$ be a vector of linearly independent forms in $\mathbf{x}$ such that $\mathbf{x}_{\text{obs}} = \mathbf{A}\mathbf{x}$. Then we say that a lumping $\mathbf{y} = \mathbf{L}\mathbf{x}$ is a constrained linear lumping with observables $\mathbf{x}_{\text{obs}}$ if each entry of $\mathbf{x}_{\text{obs}}$ is a linear combination of the entries of $\mathbf{y}$.*

2.2 The nonuniqueness/interpretability issue

A recent software CLUE [16] allows to find, for a given system (1) and a vector $\mathbf{x}_{\text{obs}}$, a constrained linear lumping of the smallest possible dimension. However, such an optimal lumping is not unique in the following sense: if $\mathbf{y}_1 = \mathbf{L}\mathbf{x}$ is a constrained linear lumping of the smallest possible dimension, then, for every invertible matrix T of the appropriate dimension, $\mathbf{y}_2 = T\mathbf{L}\mathbf{x}$ is also such a

lumping. Two such lumpings will be called *equivalent*, and one can show that all constrained linear lumpings of the smallest possible dimension are equivalent.

Because of this nonuniqueness, the lumping produced by CLUE will be optimal in terms of the dimension but not necessarily optimal in terms of the *interpretability* of the resulting macro-variables. For example, the macro-variables constructed by CLUE for the system (3) are:

$$y_1 = [X], \quad y_2 = [A_{XU}] + [A_{UX}] + 2[A_{XX}], \quad y_3 = [A_{UU}] - [A_{XX}].$$

The last macro-variable is different from the one in (4) and does not allow for the “concentration-of-sites” interpretation. Moreover, the reduced ODE system is more complicated than (5). This issue becomes more serious in more realistic (and larger) models: for the case studies in [16, Section 4.2] it has been observed that some of the resulting macro-variables “escaped physical intelligibility”.

2.3 Our approach via nonnegativity

It has been already observed in [16, Section 4.2] that the macro-variables involving negative coefficients (such as $[A_{UU}] - [A_{XX}]$) may be an obstacle for interpretability. This is partially because such quantities cannot be naturally viewed as concentrations of some sort since they may take on negative values.

Thus, in order to improve the interpretability of a lumping, we construct an equivalent lumping with all the coefficients being nonnegative and the number of nonzero coefficients (that is, the ℓ_0 -norm $\|\cdot\|_0$) being the smallest possible under the nonnegativity constraint. Mathematically, for a given lumping $\mathbf{y}_1 = L\mathbf{x}$, we find (if possible) an equivalent lumping $\mathbf{y}_2 = TL\mathbf{x}$ with invertible T satisfying:

1. the entries of TL are nonnegative and
2. $\|TL\|_0$ is as small as possible.

As we have mentioned, for fixed observables, all the constrained linear lumpings of the smallest dimension are equivalent, so the value $\|TL\|_0$ does not depend on the choice of L in the case of the optimal constrained linear lumping as in [16].

We hypothesize that the new lumping $\mathbf{y}_2 = TL\mathbf{x}$ will be typically more interpretable than the original one. We support this hypothesis by three case studies: multisite protein phosphorylation [18], Fcε-RI signaling pathways [8], and Jak-family protein tyrosine kinase activation [2]. The first two are exactly the case studies from [16] for which some of the macro-variables could not be properly interpreted by the authors.

2.4 Algorithmic details

In this section, we provide and justify Algorithm 1, an algorithm for computing a new lumping described in Section 2.3. We will use some basic terminology from convex geometry. We refer the reader to [17, Chapters 7-8] for details.

Throughout the rest of the section, for A being a vector or a matrix, $\|A\|_0$ denotes the ℓ_0 -norm of A , that is, the number of nonzero entries in A .

Algorithm 1: Algorithm for constructing new lumping

Input a $m \times n$ matrix L with entries in $\mathbb{R}$ and linearly independent rows;

Output an invertible $m \times m$ matrix T such that

- the entries of TL are nonnegative
- and the number of the nonzero entries is as small as possible.

Returns NO if such matrix T does not exist.

- (**Step 1**) Consider the row space of L and the nonnegative orthant in $(\mathbb{R}_{\geq 0})^n$ as polyhedral cones C_1 and C_2 in $\mathbb{R}^n$.
- (**Step 2**) Compute a polyhedral cone $C = C_1 \cap C_2$. This can be done, for example, using the Fourier-Motzkin algorithm [20, Section 1.2].
- (**Step 3**) If $\dim C < m$, return NO
- (**Step 4**) Let E be a set of representatives of the extreme rays of C .
- (**Step 5**) Initialize a $0 \times n$ matrix L_1
- (**Step 6**) While $E \neq \emptyset$
- (a) choose $e \in E$ such that $\|e\|_0 = \min_{v \in E} \|v\|_0$;
 - (b) if e is not in the row space of L_1 , append e to L_1 as a new row;
 - (c) remove e from E .
- (**Step 7**) Construct an $m \times m$ matrix T such that the i -th column contains the coordinates of the i -th row of L_1 with respect to the rows of L .
-

Remark 2 (Implementation). Our implementation of Algorithm 1 in Julia can be found at <https://github.com/xjzhaang/LumpingPostiviser>. We used poly-make [11] for operations with cones (at (**Step 2**) and (**Step 4**)) and Nemo [10] for symbolic linear algebra (at (**Step 6**)). Table 1 below summarizes the performance of the code on the case studies we discuss in this paper. We also provide timing for obtaining the starting reduction using CLUE. Therefore, the sum of the last two columns is the total time to obtain the final reduction for the original system. The runtimes are measured on a laptop with a 2.20GHz CPU and 16GB RAM using `@btime` macro in Julia. One can see that the models with hundreds of equations can be tackled in less than a minute on a commodity hardware.

Model	# original variables (n)	# macro-variables (m)	Runtime (sec.)	
			CLUE	Algorithm 1
Section 3.1, $m = 2$	18	6	< 0.01	< 0.01
Section 3.1, $m = 3$	66	6	< 0.01	< 0.01
Section 3.1, $m = 4$	258	6	0.34	3.4
Section 3.2	354	69	3.3	4.7
Section 3.3	470	322	72	49

Table 1: Running times of our implementation.

Remark 3 (Choice at (Step 6)a). At the (Step 6)a, if there are several $e \in E$ with $\|e\|_0$ being minimal possible, we choose the one with the index of the leftmost nonzero entry being the smallest one. In our experience, this makes the results slightly easier to analyze.

Remark 4 (Returning NO). Although Algorithm 1 may, in principle, return NO, we did not encounter such a situation with models from the literature. We give an artificial example with this property in Appendix.

Theorem 1 (Correctness of Algorithm 1). *For every matrix L over $\mathbb{R}$ with linearly independent rows, Algorithm 1 produces an invertible square matrix T such that*

- TL has nonnegative entries
- and the number of nonzero entries in TL is the smallest possible under the nonnegativity constraint

if such T exists and returns NO if there is no such T .

Proof. First, we will show that the algorithm returns NO if and only if there is no such matrix. Assume that there is such a matrix T . Then both C_1 and C_2 contain the rows of the matrix TL . Therefore, C contains m linearly independent vectors, so its dimension is at least m . In the other direction, if $\dim C \geq m$, then there exist m linearly independent vectors in $C = C_1 \cap C_2$. Let T be the matrix with the columns being their coordinates with respect to the rows of L . Then the rows of TL will belong to C_2 so that they will be nonnegative.

Now assume that the algorithm does not return NO. We observe that the entries of L_1 are nonnegative because all its rows belong to C_2 . The rows of L_1 belong to C_1 , so they are linear combinations of the rows of L . Since, by the construction on (Step 6), the rows of L_1 are linearly independent, and there are $\dim C$ of them, we conclude that the row spaces of L_1 and L coincide. Therefore, the coordinates in (Step 7) are well-defined, so the algorithm will produce a matrix T such that $L_1 = TL$ has only nonnegative entries.

It remains to prove that the ℓ_0 -norm $\|L_1\|_0$ of $L_1 = TL$ is the smallest possible. Consider any set S of m linearly independent elements of the set E of representatives of the extreme rays of C . Since (Step 7) is a greedy algorithm on the linear matroid defined by E , [7, (18)] implies that

$$\|L_1\|_0 \leq \sum_{e \in S} \|e\|_0. \quad (6)$$

Consider any invertible matrix $\tilde{T}$ such that the entries of $\tilde{L}_1 := \tilde{T}L$ are nonnegative. Since the rows $r_1, \dots, r_m$ of $\tilde{L}_1$ belong to C , each of them can be represented as a nonnegative combination of the elements of E [17, §8.8]. For each $i = 1, \dots, m$, we fix such a representation for r_i and denote $E_i \subseteq E$ the set of elements of E appearing in the representation with positive coefficients. We apply the generalized Hall's theorem [19, Theorem 1] to the family $\mathcal{A} = \{E_1, \dots, E_m\}$ of subsets of E and the function μ such that $\mu(S)$ is defined to be the dimension of the linear span of the elements of S for every $S \subseteq E$. This yields linearly

independent elements $e_1, \dots, e_m \in E$ such that $e_i \in E_i$ for every $i = 1, \dots, m$. For every $i = 1, \dots, m$, r_i is a positive combination of e_i and maybe some other elements of E , hence $\|e_i\|_0 \leq \|r_i\|_0$. Using (6), we have

$$\|L_1\|_0 \leq \sum_{i=1}^m \|e_i\|_0 \leq \sum_{i=1}^m \|r_i\|_0 = \|\tilde{L}_1\|_0,$$

and this proves the minimality of the number of the nonzero entries in $L_1 = TL$ for T constructed by the algorithm.

3 Case Studies

In this section, we demonstrate the improvements in physical intelligibility (while preserving the dimension) of reductions of biochemical models by our Algorithm 1. We analyse the results of the algorithm using models taken from the literature. We also compare the resulting reduction to the ones obtained by ERODE [4] which are always defined by zero-one linear combinations.

3.1 Multisite protein phosphorylation

Setup. We consider a model of multisite phosphorylation [18]. It describes a protein with m identical and independent binding sites that simultaneously undergo phosphorylation and dephosphorylation. Each binding site can be in one of the four different states (see Figure 1):

1. unphosphorylated and unbound,
2. unphosphorylated and bound to a kinase,
3. phosphorylated and unbound,
4. phosphorylated and bound to a phosphatase.

Therefore, there are $4^m + 2$ chemical species in the corresponding reaction network: free kinase and phosphatase, and 4^m states of the protein.




Fig. 1: Molecular components and states of the multisite phosphorylation model.
a. Consists of multisite proteins and kinases. This example has $m = 3$ sites.
b. There are 4 possible states for a single site: unphosphorylated and unbound, unphosphorylated and bound to a kinase, phosphorylated and unbound, phosphorylated and bound to a phosphatase.

Reductions by ERODE and CLUE. In the reduction computed by ERODE [5] (for $m = 2, \dots, 8$), the concentrations of protein configurations are replaced by the sums of the concentrations of configurations differing by a permutation of the sites. Therefore, the number of macro-variables is equal to $\binom{m+3}{3} + 2$.

In contrast, the analysis performed by CLUE [16] always results in just six macro-variables. Two of them were always the concentrations of kinase and phosphatase as for ERODE. The other four were linear combinations with protein configurations. In [16, Section 4.2], for $m = 2$, interpretation was provided for the first three of them. However, for the last one, it was remarked that “the last macro-variable escaped physical intelligibility as it represents the difference between the free substrate with unphosphorylated sites and protein configurations that appear in the aforementioned lumps.”

Our results. We applied our algorithm to the cases $m = 2, 3, 4, 5$ and obtained new macro-variables, which have again included the concentrations of free kinase and phosphatase. Moreover, the three interpretable macro-variables from the analysis in [16] for $m = 2$ are kept. Each of the four our macro-variables involving the protein configurations corresponds to a state of a site (e.g., unbounded and unphosphorylated), and each protein configuration appears with a coefficient equal to the number of sites in it with this state. Examples of these new macro-variables are given on Figure 2 for $m = 2$ and $m = 3$.

One way to interpret the result is that the constructed reduction replaces the concentration of the protein configurations with the “concentrations” of each of

the four states of the sites (see also Example 1). From our interpretation, we expect that the models with larger m will have a reduction of the same form.




Fig. 2: New macro-variables for $m = 2, 3$. Each state of a binding site from Fig 1-b can be the main state, yielding four macro-variables for each m . The coefficients are equal to the number of binding sites in a protein that are in the main state.

3.2 Fcε-RI signaling pathways

Setup. We consider a model for a different kind of multisite phosphorylation [8], a model for the early events in the signaling pathway of the high-affinity IgE receptor (FcεRI) in mast cells and basophils.

The model details the rule-based interactions of FcεRI receptor with a bivalent ligand (IgE dimer), the Src kinase Lyn, and the cytosolic protein tyrosine kinase Syk. The model is based on the following sequence of signaling events in FcεRI [13,14] (the reactions are nicely summarized on [8, Figure 2]):

1. binding of IgE ligand and FcεRI which aggregates at the plasma membrane,
2. transphosphorylation of tyrosine residues in the immunoreceptor tyrosine-based activation motifs (ITAMs) of the aggregated receptor by constitutively associated Lyn,
3. recruitment of an extra Lyn/Syk kinase to the phosphorylated ITAM sites,
4. transphosphorylation of Syk by Lyn and Syk on its linker region and activation loop, respectively.

For visualizing different chemical species occurring in the resulting reaction network, we use the representation [8, Figure 1] summarized in Figure 3. In total, there are 354 of three types: monomers, dimers, and non-receptor states (free ligand/Lyn and Syk in each of 4 possible states of phosphorylation).




Fig. 3: Molecular components and states of the FcεRI signaling events model.
a. IgE dimer is a bivalent ligand. FcεRI consists of α , β , γ subunits. Lyn kinase has an SH2 domain. Syk kinase has an SH2 domain and two ITAM sites which differ by the method of phosphorylation: Lyn at the linker region, and Syk at the activation loop.
b. The α subunit can be unbound or bound to a ligand. β can be unphosphorylated/phosphorylated, with/without associated Lyn. γ can be unphosphorylated/phosphorylated, and the phosphorylated form binds to Syk in any of the four states of phosphorylation.




Fig. 4: Examples of a monomer, a dimer, and the free components.

Reductions by ERODE and CLUE. The reduction by ERODE [5] consists of 105 macro-variables, where all the complexes with the same configuration except for the phosphorylation state of the Syk units are summed up in a single macro-variable. We will refer to these macro-variables as *Syk-macro-variables*.

The model has been reduced using CLUE in [16, Section 4.2] with the observable being the total concentration of the free Syk (in all the four phosphorylation states). The reduced model had 69 macro-variables, and 51 of them were nontrivial. It has been observed in [16, Section 4.2] that some of these macro-variable carry a physical interpretation, but in some of them, negative elements were present, which hinder their physical intelligibility.

Our results. We apply our algorithm to the reduced model. Among the new macro-variables, we have 51 nontrivial macro-variables which is the same as for the CLUE reduction. More precisely, **(Step 4)** produced 57 nontrivial macro-variables, and this number has been reduced to 51 when computing a linearly independent basis on **(Step 6)**. The resulting macro-variable refine the reduction by ERODE mentioned above in the sense that our new macro-variable are the sums of the Syk-macro-variables with non-negative coefficients. Therefore, in our reduction, all the complexes differing only by the phosphorylation state of the Syk units are in the same macro-variable. For the *monomers*, we obtain the same reduction: nontrivial macro-variable involving monomers are of the form described on Figure 5.




Fig. 5: The monomer macro-variables. In each of them, α and β are fixed, and we sum over all the phosphorylation patterns of Syk.

The macro-variable involving dimers are graphically described on Figure 6. First, one can see that they are indeed linear combinations of the Syk-macro-variables. Our interpretation of these new macro-variables is based on two observations about the set of the reactions in the original model [8, Figure 2]:

- (Obs. 1) For every reaction involving a dimer, only one of the receptors of the dimer is affected by the reaction.
- (Obs. 2) The γ -chain of the other (not affected) receptor is relevant only for the reactions of transphosphorylation of Syk.


Since the complexes with different phosphorylation patterns are grouped together in the Syk-macro-variables, the second observation implies that the transphosphorylation reactions do not affect the values of the Syk-macro-variables at all. Therefore, the first observation suggests considering macro-variables as sums over all the dimer configurations in which one receptor is fixed (up to the phosphorylation of Syk), and for the other receptor, all the possible variants of the γ -chain are considered.

With this interpretation in mind, let us take a closer look at the Figure 6:


- Each of the variables as on Figure 6a is the sum over the configurations with the fixed left receptor not carrying Syk and the right receptor having each of the six possible γ -chains. If one of the complexes in the sum is fully symmetric, it appears with coefficient 2.
- Each of the variables as on Figure 6b is a combination of complexes that have: the same β -chains and Syk on the left receptor, any phosphorylation pattern

of the Syk on the left receptor, and any γ -chain on the right receptor. If the β -chains on the receptors are equal, the complexes with two Syk's (which are *symmetric up to Syk phosphorylation*) appear with coefficient 2.


Note that the coefficients 2 appearing in the presence of symmetry prevent ERODE [5] from finding this reduction.



(a) No Syk on the left γ -chain, the right takes all the possible γ -configurations



(b) Syk (with all possible phosphorylations) on the left γ -chain, the right takes all the possible γ -configurations.



(c) An example configuration of Y_3 (b) with identical β -chains. A variable has coefficient 2 when Syk kinases are recruited on both receptors (regardless of Syk state).

Fig. 6: Macro-variables involving dimers

3.3 Jak-family protein tyrosine kinase activation

Setup. We study a simplified cellular model of a bipolar “clamp” mechanism for Jak-family kinase activation [2]. Kinases of the Janus kinase (JAK) family play an essential role in signal transduction mediated by cell surface receptors, which lack innate enzymatic activities to dimerize.

The model studies the interactions of Jak2 kinase trans-phosphorylation, specifically the rule-based dynamics between the Jak2 (J) kinase, the unique adaptor protein SH2-B β (S) with the capacity to homo-dimerize, the growth hormone receptors (R), and a bivalent growth hormone ligand (L). The Jak2 kinase has two phosphorylation sites, Y1 and Y2. The SH2-B β protein contains an N-terminal dimerization domain (DD) and a C-terminal Src homology-2 (SH2) domain.

The components can interact in the following ways:

1. binding of ligand and growth hormone receptors which aggregates at the plasma membrane,
2. constitutive binding of Jak2 kinase to the receptors, which autophosphorylates on the phosphorylation sites when two Jak2 kinases are bounded in the same complex,
3. recruitment of SH2-B β protein at the SH2 domain by the Jak2's autophosphorylated Y1 site,
4. dimerization of SH2-B β protein through recruitment of an additional SH2-B β protein, engaged at the DD domains.

Receptors can undergo a process of *internalization*, in which the receptors can no longer associate with any Jak2, and the existing Jak2 and SH2-B β in the complex can dissociate at the normal rate.




Fig. 7: Components of the Jak-family protein tyrosine kinase activation model.

Reductions by ERODE and CLUE. The reduction obtained by ERODE in [5, Figure 5] contained 345 macro-variables. It grouped the internalized configurations, which differ by the connections between the receptors and the ligand, into macro-variables.

The model has been reduced by CLUE in [16], with the observable being the concentration of the free ligand. The reduced model had 322 macro-variables, and 69 of them were nontrivial. The model has been used in [16] for benchmarking purposes only, so the macro-variables have not been interpreted. The reduction included several macro-variables with negative coefficients, including one with 20 nonzero coefficients. We do not see any natural interpretation for them.

Our results. We apply our algorithm to the reduced model, and among the produced macro-variables 69 are nontrivial as in the reduction by CLUE ((**Step 4**) produces 70 macro-variables, and then this number is reduced to 69 at (**Step 6**)). The nontrivial macro-variables are linear combinations of internalized molecules, and the trivial macro-variables are not internalized. Compared to the ERODE

reduction, some internalized complexes such as the ligand-receptor (R, RL, RLR) structures in [5, Figure 5(C)], are omitted in our model as they do not disassociate under internalization and thus do not affect the dynamics of the free ligand observable. In our reduction, mirrored internalized complexes are lumped together, which explains all two-element macro-variables. The remaining non-trivial macro-variables are described on Figure 8, and are of two types:

- *Configurations equivalent up to the connection between the ligand and the receptors (Figure 8-a and 9).* The structures are equivalent under internalization as the ligand and receptors cannot disassociate and were obtained also by ERODE [5, Figure 5(D, E)]. They are of two types: single-Jak2-Receptor case (Figure 8-a) and “clamp” case (Figure 9).
- *Configurations with one receptor fixed (Figure 8-b).* These are similar to Figure 6 from the case study in Section 3.2: since in the reactions with internalized complexes, only one receptor is affected, and this does not depend on the state of the other receptor, one can group together the complexes having one of the receptors the same.

When the receptors are symmetric, the element and its mirrored element are the same, so the corresponding configuration appears with coefficient 2.




Fig. 8: Classes of internalized macro-variables. a. equivalent up to the connection between the ligand and the receptors. b. one receptor fixed.




Fig. 9: Macro-variables for equivalent structures of the bipolar “clamp” mechanism

4 Conclusion

We have hypothesized that the interpretability of the macro-variables in an exact linear reduction may be improved by a change of coordinates making the macro-variable nonnegative combinations of the original variables and minimizing the number of nonzero coefficients. We have designed and implemented an algorithm for performing such a transformation and applied it to three models (with hundreds of variables) for which the result of the reduction by CLUE [16] contained macro-variables without a clear physical interpretation. We have shown that the resulting macro-variables are interpretable, thus supporting the original hypothesis and demonstrating the usefulness of our algorithm.

Our results also give insight into the structure of reductions in which not all the coefficients are zeroes and ones. In particular, we can point out two different situations:

- The macro-variables are the “concentrations” of parts of molecules as in Section 3.1. The species having several identical pieces may appear with larger coefficients.
- Some of the molecules appearing in the macro-variable are symmetric (as in Section 3.3) or even partially symmetric (as in Section 3.2), and they appear with a coefficient accounting for the symmetries.

Acknowledgement

The authors are grateful to François Fages, Mathieu Hemery, Sylvain Soliman, and Mirco Tribastone for helpful discussions and to the referees for helpful suggestions. GP is grateful to Heather Harrington, Gregory Henselman-Petrusek, and Zvi Rosen for educating him about the matroid theory.

References

1. Antoulas, A.: Approximation of Large-Scale Dynamical Systems. Adv. in Design and Control, SIAM (2005)
2. Barua, D., Faeder, J.R., Haugh, J.M.: A bipolar clamp mechanism for activation of Jak-family protein tyrosine kinases. PLoS Computational Biology **5**(4), e1000364 (2009), <https://doi.org/10.1371/journal.pcbi.1000364>
3. Borisov, N., Markevich, N., Hoek, J., Kholodenko, B.: Signaling through receptors and scaffolds: Independent interactions reduce combinatorial complexity. Biophysical Journal **89**(2), 951–966 (2005), <http://dx.doi.org/10.1529/biophysj.105.060533>
4. Cardelli, L., Tribastone, M., Tschaikowski, M., Vandin, A.: ERODE: A tool for the evaluation and reduction of ordinary differential equations. In: TACAS 2017. LNCS, vol. 10206 (2017), https://doi.org/10.1007/978-3-662-54580-5_19
5. Cardelli, L., Tribastone, M., Tschaikowski, M., Vandin, A.: Maximal aggregation of polynomial dynamical systems. Proceedings of the National Academy of Sciences **114**(38), 10029–10034 (2017), <https://doi.org/10.1073/pnas.1702697114>

6. Conzelmann, H., Fey, D., Gilles, E.: Exact model reduction of combinatorial reaction networks. *BMC Systems Biology* **2**(1), 78 (2008), <http://dx.doi.org/10.1186/1752-0509-2-78>
7. Edmonds, J.: Matroids and the greedy algorithm. *Mathematical Programming* **1**(1), 127–136 (1971), <https://doi.org/10.1007/bf01584082>
8. Faeder, J.R., Hlavacek, W.S., Reischl, I., Blinov, M.L., Metzger, H., Redondo, A., Wofsy, C., Goldstein, B.: Investigation of early events in FcεRI-mediated signaling using a detailed mathematical model. *The Journal of Immunology* **170**(7), 3769–3781 (2003), <https://doi.org/10.4049/jimmunol.170.7.3769>
9. Feret, J., Danos, V., Krivine, J., Harmer, R., Fontana, W.: Internal coarse-graining of molecular systems. *Proceedings of the National Academy of Sciences* **106**(16), 6453–6458 (2009), <http://dx.doi.org/10.1073/pnas.0809908106>
10. Fieker, C., Hart, W., Hofmann, T., Johansson, F.: Nemo/Hecke: Computer algebra and number theory packages for the Julia programming language. In: *Proceedings of the 2017 ACM on International Symposium on Symbolic and Algebraic Computation*. pp. 157–164. ISSAC '17, ACM, New York, NY, USA (2017). <https://doi.org/10.1145/3087604.3087611>, <http://doi.acm.org/10.1145/3087604.3087611>
11. Gawrilow, E., Joswig, M.: polymake: a framework for analyzing convex polytopes. In: *Polytopes - Combinatorics and Computation*, pp. 43–73. Birkhäuser Basel (2000), https://doi.org/10.1007/978-3-0348-8438-9_2
12. Gunawardena, J.: Multisite protein phosphorylation makes a good threshold but can be a poor switch. *Proceedings of the National Academy of Sciences* **102**(41), 14617–14622 (2005), <https://doi.org/10.1073/pnas.0507322102>
13. Metzger, H., Eglite, S., Haleem-Smith, H., Reischl, I., Torigoe, C.: Quantitative aspects of signal transduction by the receptor with high affinity for IgE. *Molecular Immunology* **38**(16-18), 1207–1211 (2002), [https://doi.org/10.1016/s0161-5890\(02\)00065-2](https://doi.org/10.1016/s0161-5890(02)00065-2)
14. Nadler, M.J., Matthews, S.A., Turner, H., Kinet, J.P.: Signal transduction by the high-affinity immunoglobulin E receptor FcεRI: Coupling form to function. *Advances in Immunology*, vol. 76, pp. 325–355. Academic Press (2001). [https://doi.org/https://doi.org/10.1016/S0065-2776\(01\)76022-1](https://doi.org/https://doi.org/10.1016/S0065-2776(01)76022-1), <https://www.sciencedirect.com/science/article/pii/S0065277601760221>
15. Okino, M., Mavrovouniotis, M.: Simplification of mathematical models of chemical reaction systems. *Chemical Reviews* **2**(98), 391–408 (1998), <https://doi.org/10.1021/cr9502231>
16. Ovchinnikov, A., Verona, I.P., Pogudin, G., Tribastone, M.: CLUE: exact maximal reduction of kinetic models by constrained lumping of differential equations. *Bioinformatics* (2021), <https://doi.org/10.1093/bioinformatics/btab010>
17. Schrijver, A.: *Theory of Linear and Integer Programming*. Wiley–Blackwell (1986)
18. Sneddon, M.W., Faeder, J.R., Emonet, T.: Efficient modeling, simulation and coarse-graining of biological complexity with NFsim. *Nature Methods* **8**(2), 177–183 (2010), <https://doi.org/10.1038/nmeth.1546>
19. Welsh, D.: Generalized versions of Hall’s theorem. *Journal of Combinatorial Theory, Series B* **10**(2), 95–101 (1971), [https://doi.org/10.1016/0095-8956\(71\)90069-4](https://doi.org/10.1016/0095-8956(71)90069-4)
20. Ziegler, G.M.: *Lectures on Polytopes*. Springer-Verlag New York (1995), <https://doi.org/10.1007/978-1-4613-8431-1>

Appendix: Non-positivizable reduction

As we have mentioned in Remark 4, we did not encounter examples from the literature for which Algorithm 1 would return NO. However, one can easily construct an artificial example with this property. Consider the system

$$\begin{cases} x_1' = x_1^2 + x_2^2, \\ x_2' = 2x_1x_2. \end{cases} \quad (7)$$

Then $y = x_1 - x_2$ yields a reduced system $y' = y^2$. However, since any change of macro-variables is a scaling of y , there is no equivalent lumping with nonnegative coefficients, so Algorithm 1 (with the input $L = (1 \ -1)$) will return NO.