Authors' Response to Anonymous Referee #1

GENOA v3: A flexible framework for reduction and exploration of highly detailed chemical mechanisms

Zhizhao Wang, William P. L. Carter, Julia Lee-Taylor, John Orlando, Qing Ye, Richard Valorso, Marie Camredon, Bernard Aumont, and Kelley Barsanti

Geoscientific Model Development, https://doi.org/10.5194/egusphere-2026-1921

Referee comment DOI: https://doi.org/10.5194/egusphere-2026-1921-RC1

RC: Referee's Comment, AR: Authors' Response, ☐ Manuscript Text

General Comments

RC: The paper presents the GENOA v3 framework for reduction of atmospheric chemical mechanisms with huge number of species and reactions such as the highly detailed VOC reaction mechanisms generated by GECKO-A. The work demonstrates that GENOA v3 can effectively reduce highly detailed VOC mechanisms while preserving SOA mass production and gas-phase reactivity with reasonable accuracy across a wide range of atmospherically relevant scenarios.

AR: We thank the referee for the careful review and detailed, constructive comments. We have carefully considered each comment and revised the manuscript accordingly. The comments helped clarify the respective roles of TBR and SBR and raised important questions regarding SBR error control, the reproducibility of the reduction process, and the reporting of reduced mechanisms.

The responses below address each general and specific comment individually. The line and section numbers cited in the responses refer to the manuscript before revision.

RC: The reduction follows a rule-based strategy and is performed in two subsequent approaches, first through a threshold-based reduction (TBR) and second through a simulation-based reduction (SBR) that uses the TBR-reduced mechanism as a starting point. Whereas TBR provides a rapid initial pruning of large mechanisms at low computational demand, only the subsequent SBR provides the aggressive reduction needed for application in a large-scale CTM. It is not apparent whether the TBR part can be skipped in principle or whether it is required for reordering or attenuating the starting-point mechanism before it can be used in SBR. The authors currently imply that the starting mechanism needs some refinement before it is possible to run it efficiently in SBR. However, GECKO-A should already generate a mechanism that is not too stiff for theoretically applying it in a 0-D box model.

AR: We apologize for the ambiguity. TBR is not a formal prerequisite for SBR and can be skipped. SBR can begin with any chemical mechanism that can be processed by GENOA v3. We adopted the TBR–SBR workflow because the computational cost of SBR strongly depends on the size of the input mechanism. Applying TBR first can therefore substantially shorten the SBR reduction time. We have clarified this at the beginning of the Sect. 3 in the revised manuscript:

Sect. 3, page 6, line 135

~~As illustrated in Fig. 1, reductions begin from a starting-point mechanism and are evaluated relative to a reference mechanism, which may be identical in some applications. The reduction framework consists of~~ The mechanism reduction framework supports sequential and iterative reduction through two main processes~~:~~ , described below, with mechanism performance evaluated against predefined reference mechanisms:

- Threshold-based reduction (TBR) performs system-level reduction using strategies and threshold values. In each TBR step, a single reduction strategy is applied consistently across the entire mechanism. This allows rapid reduction of reactions and species and produces a substantially smaller mechanism ~~that is suitable for further refinement~~for subsequent reduction or application.
- Simulation-based reduction (SBR) performs stepwise reduction guided by simulation feedback. ~~Starting from TBR-reduced mechanisms,~~ SBR proceeds through a sequence of stages with stage-specific accuracy constraints. Within each stage, reduction is applied iteratively to sub mechanisms, with multiple reduction candidates evaluated in parallel using metrics from 0-D simulations. This allows further accuracy-controlled size reduction and produces compact mechanisms manageable for ~~use in~~ large-scale ~~models such as 3-D chemical transport models.~~CTM applications.

~~These two processes are modular but typically used sequentially in GENOA v3, with TBR followed by SBR. This design allows computationally efficient pruning of large mechanisms at the system level before applying the more computationally demanding SBR for greater size reduction with user-controlled accuracy.~~ TBR and SBR are modular and can be applied independently. In the workflow used in this study and illustrated in Fig. 1, TBR is applied first to reduce the mechanism size efficiently and thereby lower the computational cost of the subsequent SBR, which provides further reduction under user-defined accuracy constraints.

This point is also addressed again in the discussion:

Sect. 7.1, page 46, line 946

~~The presented examples illustrate these methodological differences~~Either TBR or SBR can be applied independently to any starting-point mechanism compatible with GENOA v3, although they are used sequentially in this study. These examples also demonstrate their different reduction capabilities and computational demands. For U8, TBR (U8StP) retained 20.7 % of the original size, while ~~the SBR decreased it further~~ subsequent SBR reduced the mechanism further to below 0.05 % (U8SOA-H and U8multi-H). Such reductions cannot be achieved by TBR alone while preserving reasonable accuracy. ~~However~~At the same time, the high computational demand for SBR limits its direct application to large VOC mechanisms. ~~As a result, a combined workflow that applies TBR first and then SBR provides a reasonable and scalable approach, enabling the preservation of accuracy while achieving extreme size reductions across diverse mechanisms and~~ Applying TBR before SBR therefore provides a practical and scalable workflow: TBR lowers the computational cost of SBR, while SBR enables substantial further reduction under user-defined ~~objectives~~accuracy constraints.

To further assess the computational benefit of applying TBR before SBR, we performed an additional SOA-focused U8 reduction in which SBR was initialized directly from Ref using the same Stage A1 error tolerances

as in the standard workflow. The results have been added to Sect. S3.1.3 of the Supplement. Direct SBR successfully reduced Ref to A1′, containing 646 reactions, 89 species, and 34 condensable species, with mean and maximum errors of 4.9 and 5.8 %, respectively. However, this required 21.93 kCPU-h, compared with 3.59 kCPU-h for TBR followed by Stage A1. Most of this difference arose from the first SBR cycle, which required 21.56 kCPU-h when initialized from Ref, compared with 3.3 kCPU-h when initialized from StP. Thus, SBR can be applied directly without TBR, but prior TBR substantially reduces the computational cost of the most demanding initial SBR cycle. This result is referenced in Sect. 7.2.3 as follows:

> *Sect. 7.2.3, page 48, line 1000*
>
> ~~More aggressive TBR prior to SBR , or the introduction of an~~ Using the same Stage A1 settings, SBR initialized directly from the U8 Ref mechanism required 21.93 kCPU-h to complete Stage A1, compared with 3.59 kCPU-h for the TBR–SBR workflow (Table S8; Sect. S3.1.3). Prior TBR, or a similar intermediate reduction stage, may ~~help address this limitation~~ therefore help reduce the computational cost of SBR for very large mechanisms.

RC: SBR is not a black box but the reduction comes with the risk that the reduced mechanism does not remain accurate under all given or even not the most critical atmospheric conditions.

AR: We agree. The accuracy of an SBR-reduced mechanism depends on the evaluation targets, atmospheric scenarios, tolerance settings, and other reduction configurations. Its performance can therefore only be established for the conditions represented by the evaluation scenarios, and its accuracy under conditions encountered in 3-D applications requires further evaluation. This highlights the need to tailor scenario sets and evaluation targets to the intended application. This scenario dependence is clarified in Sect. 7.2.1:

> *Sect. 7.2.1, page 47, line 971*
>
> Consequently, reduced mechanisms are optimized for conditions represented by the selected scenarios, and alternative or expanded scenario sets are likely to produce different reduction outcomes. Constructing representative scenario sets remains challenging because atmospheric conditions are highly diverse and complex, making it difficult to extract the relevant chemical and physical information into a limited number of representative scenarios . . .

RC: The error control during mechanism reduction is somewhat vague. Are there any checks and controls happening during the reduction process chain such as atom balance (loss of C or N or O)?

AR: Thank you for raising this point. SBR does not currently impose a global elemental-balance criterion for C, N, or O. However, molecular composition is considered when identifying candidates for some strategies: lumping constrains differences in carbon and oxygen atom numbers among lumpable species, while replacement constrains differences in carbon number together with the structural and kinetic similarity criteria listed in Table B3. These criteria limit compositional differences among reducible species, but they do not ensure exact elemental conservation. No explicit elemental-balance criterion is currently applied to removal or jumping candidates.

Because SBR candidate acceptance is based on the simulation-error criteria rather than an elemental-budget test, elemental loss may therefore occur during reduction. An elemental-budget metric could be introduced as an additional acceptance criterion in future applications. The elements to constrain and their acceptable deviations would need to be defined according to the intended application and require further evaluation. We have revised Sect. 7.2.2 accordingly:

Sect. 7.2.2, page 47, line 991

7.2.2 Configuration sensitivity and ~~carbon~~ elemental-budget conservation

. . .

In the current implementation, the reduction workflow does not enforce global elemental conservation (e.g., of C, O, or N) and tends to favor removal-based strategies, which can lead to increased carbon loss in highly reduced mechanisms. This behavior reflects an intentional trade-off between mechanism compactness and strict carbon conservation when removals are prioritized. A diagnostic analysis quantifying carbon propagation across reduction configurations is provided in the Supplement (Table S16; Sect. S3.~~5~~.4). While such carbon loss ~~is acceptable for target-focused applications , future~~ may be acceptable for applications focused on specific evaluation targets, its suitability should be assessed for the intended use. Future developments could explore ~~placing greater emphasis on~~ lumping-based strategies with explicit carbon control~~.~~, redistribution of carbon from truncated pathways to retained products, or application-specific elemental constraints. The appropriate elements, redistribution rules, and allowable deviations would require further evaluation.

RC: The user-defined error tolerances for the reduction errors in SBR are also part of this concern. Not only does the choice of the error tolerance (ϵ) increase the degrees of freedom of a reduced mechanism created from the same starting mechanism by different users, it is also not demonstrated in this work whether the error of the reduced mechanism is related to the value of ϵ. Therefore, it is not guaranteed that decreasing ϵ results in a monotonically decreasing error of simulated concentrations compared to the reference, and may not always provide the smallest reduced mechanism at a given simulation error. Probably different epsilon values should be tested by the user and could result in a reduced mechanism with less species but only slightly different simulation error in comparison to the reference mechanism. My suggestion is to study how number of species and simulation error (NMB) change as function of epsilon for a subset of the mechanism, for example U8.

AR: Thank you for pointing out this important issue and for suggesting an additional sensitivity analysis.

We agree that user-defined error tolerances introduce degrees of freedom and that a smaller tolerance does not necessarily produce a monotonically smaller final error or the smallest possible mechanism at a given error level. SBR uses three complementary tolerances for fractional-absolute-error-type (FAE-type) reduction errors (Sect. 3.3.3; Eqs. (C4)–(C5)). The mean-error tolerance, ϵ_{avg}, constrains the mean error across evaluation targets and scenarios; $\epsilon_{\max}$ constrains the largest target–scenario error; and ϵ_{Δ} limits the increase in the mean or maximum error relative to the previously accepted mechanism. Among candidates satisfying all applicable tolerances and stage-specific restrictions, the candidate with the highest reduction score is selected.

In this study, the tolerances are progressively relaxed as reduction proceeds (Table 5), providing a practical means of controlling the size–accuracy trade-off. For U8, the representative SOA-focused mechanisms obtained with low-, medium-, and high-tolerance settings have mean errors of 4.9, 9.6, and 14.6 %, respectively, while the corresponding multi-target mechanisms have mean errors of 5.0, 9.9, and 15.0 % (Table 6). These results show that the adopted tolerance sequence provides useful control over error growth. NMB is reported as a diagnostic of systematic bias but is not used for candidate acceptance, which is based on FAE-type reduction errors.

The lack of a guaranteed monotonic relationship noted by the reviewer is an important consequence of the path dependence of SBR. In SBR, error tolerances define acceptance bounds for individual reduction steps, but they do not uniquely determine the final mechanism. Each accepted reduction changes both the mechanism and the candidates evaluated in later steps, so different tolerance sequences or stage configurations

can lead to different reduction pathways. This path dependence is further clarified in Sect. 7.2.3 of the revised manuscript:

Sect. 7.2.3, page 47, line 986

7.2.3 ~~Computational cost~~Path dependence, computational cost, and scalability of SBR

SBR is a path-dependent reduction process. Each accepted reduction changes the mechanism and the candidates evaluated in subsequent steps. Consequently, different error tolerances, stage configurations, or small numerical differences for candidates close to an acceptance threshold can lead to different reduction decisions and final mechanisms. SBR therefore does not guarantee a unique or globally smallest mechanism for a prescribed error level. An exhaustive search over all possible reduction pathways would be computationally impractical for the mechanisms considered here because the number of candidate combinations grows rapidly with mechanism size. For smaller mechanisms or later reduction stages, however, it may be feasible to retain several acceptable candidates and explore multiple pathways, producing a set of mechanisms with different size–accuracy trade-offs rather than a single result from one greedy pathway. This option is not implemented in the present study, but could be explored in future GENOA applications.

Following the reviewer's suggestion, we also added a U8 sensitivity test of the SBR error-tolerance settings in a new Sect. S3.1.4 and Table S9 of the Supplement, where the analysis and results are presented in detail. Starting from U8A1 and U8B1, each final low-, medium-, and high-tolerance set was applied independently, without the preceding progressively relaxed tolerance levels used in the staged reductions reported in Table 5. The results show that the tolerance levels provide practical control over the allowable error range and the overall size–accuracy trend, but that mechanisms obtained with the same final tolerance set can differ in size and composition depending on the reduction sequence and other stage settings.

Because GENOA v3 is intended to support diverse mechanisms and applications, this limited test does not establish a universal or generally optimal tolerance strategy. The staged configurations used in this work are therefore regarded as reproducible starting points rather than settings guaranteed to outperform alternative approaches, and further investigation is needed for different mechanisms and applications. The new sensitivity test is also referenced in the main paper when describing the staged configuration and error-tolerance settings:

Sect. 4.4.2, page 22, line 514

The staged tolerance sequences were selected empirically and do not represent the only viable configuration. An additional U8 sensitivity test, in which each final low-, medium-, or high-tolerance set was applied independently starting from U8A1 or U8B1, is presented in Sect. S3.1.4 of the Supplement.

RC: The more compact chemical mechanism resulting from GENOA v3 may be viable for large-scale atmospheric simulations, although this is not demonstrated in the paper. Since each user-defined set of configurations of targets and error tolerances results in a differently reduced mechanism, it will be difficult for those users who apply the reduced mechanism in a CTM simulation to define the mechanism. The use of the GENOA-reduced mechanism in CTM simulations would require the publication of the full reduced mechanism including the user-specific reduction protocol either before or along with the CTM simulation when publishing model results. A nomenclature for the obtained reduced chemical mechanism should be suggested in this work.

AR: The use of GENOA v3-reduced mechanisms in large-scale CTM simulations has not yet been demonstrated. Previous GENOA v2 reductions of monoterpene and sesquiterpene degradation schemes from MCM (with

PRAM for autoxidation), have been applied to regional CTM simulations over Europe (Wang et al., 2024). Similar applications and evaluations remain future work for GENOA v3.

As suitable reduction settings depend on the intended application, we do not suggest a universal protocol at this stage. Instead, the GENOA v3 test cases provide example sets of reduction options in the GitHub repository (`https://github.com/tool-genoa/genoa`), with additional descriptions available on the GENOA website (`https://tool-genoa.github.io/genoa-web/test_cases/`). These test cases, together with the settings used in this study, serve as reproducible examples and starting points for future applications. Reduced mechanisms used in CTM simulations should similarly be published with their reduction settings and evaluation results.

For nomenclature, the labels used in this study indicate the precursor, evaluation target, and tolerance level, as in U8SOA-H and U8multi-H. A more general name could also include the scenario set, for example, `U8_SOA_10_urban` for a 1-octene mechanism reduced for SOA with a 10 % mean reduction error tolerance using an urban scenario set. However, a short name cannot include all reduction settings, so the full protocol should be reported separately. We do not propose a fixed nomenclature in this first GENOA v3 study, as a useful standard will require broader use of the framework and input from users across different applications.

We have clarified the need for future CTM development and evaluation, as well as the role of community use in developing practical guidance, in revised Sect. 7.3:

> *Sect. 7.3, page 48, line 1010*
>
> Future 3-D applications will require further development and evaluation of GENOA v3 reduced mechanisms for use in large-scale CTMs. Beyond individual case studies, one potential direction is the development of predefined sets of reduced mechanisms optimized for representative environments, ...
>
> ... GENOA v3 provides a platform for testing new strategies, sharing reduced mechanisms, and incorporating feedback from diverse modeling contexts. Broader use of the framework will help develop practical guidance for selecting scenarios, evaluation targets, and other settings across mechanisms and applications. Continued development will emphasize extensibility, transparency, and reproducibility, allowing the framework to evolve alongside community needs and advances in atmospheric chemistry modeling.

RC: The evaluation of the SBR reduction performance (section 6) is well organized and presented. The reaction pathway visualization is excellent and the example of reaction pathways for U8 with SOA-focused reductions at high and medium error tolerances versus multi-target reduction is well chosen. SBR is the main development of this work, and this part should be more in focus of the paper. Given the minor role of TBR in reducing the original mechanism, it would be sufficient to describe TBR in form of a technical manual in the Supplement or at least to shorten the TBR evaluation part in the manuscript.

AR: Thank you for this suggestion. SBR provides the more substantial reduction needed to produce highly compact mechanisms for 3-D CTMs. However, TBR still plays an important role in the presented workflow: it reduces the reference mechanisms by approximately 46–80 %, lowers the computational cost of subsequent SBR, and defines the starting-point mechanisms used in the SBR evaluation. Retaining its evaluation is therefore useful for quantifying the initial size–accuracy trade-off and interpreting the subsequent SBR results.

After careful consideration, we have retained the TBR evaluation in the main manuscript, while substantially shortening the methodological description in Sect. 3.2. The more technical TBR formulations and calculations, their relation to GECKO-A mechanism generation, and a newly added worked TBR demonstration using a simple example mechanism have been moved to the new Sect. S2.1 of the Supplement. This provides a more

practical introduction to the TBR procedure while keeping the main text focused. As clarified in the previous response and the corresponding revisions to Sects. 3 and 7.1, TBR and SBR can be applied independently, while their sequential use provides an efficient and practical reduction workflow.

Sect. 3.2 now provides a more concise description and directs readers to these supporting details:

> *Sect. 3.2, page 9, line 210*
>
> ~~Table B1 provides an overview~~ Additional TBR details are provided in Sect. S2.1, including mathematical formulations and related calculations for the available strategies ~~and their mathematical formulations.~~ (Table S3), the effects of prior GECKO-A generation settings on TBR, and a demonstration of TBR using a simple mechanism (Fig. S7).

RC: The development of GENOA v3 is a very useful instrument for atmospheric chemistry modeling. The evaluation of reduction candidates in terms of SOA and multiple gas-phase reactivity metrices and the capacity to reduce starting mechanisms containing hundred thousands of reactions into compact mechanisms makes this work a valuable contribution.
The revision of the manuscript should include a comparison of the GENOA-reduced mechanism against other condensed chemistry mechanisms such as the Common Representative Intermediate (CRI) scheme from reducing MCM. In addition, my concerns listed below should be addressed.

AR: Thank you for the positive assessment of GENOA v3 and for the suggestion. A fully like-for-like comparison with other condensed mechanisms would require substantial additional evaluation under harmonized model implementations and conditions and is beyond the scope of the present work, given differences in parent chemistry, chemical scope, reduction approach, and evaluation objectives. However, we agree that a mechanism-size comparison with CRI v2.2 provides a useful external reference. Among the three SBR precursors, only n-dodecane has a precursor-matched CRI v2.2 scheme. We therefore added a comparison between the CRI v2.2 NC12H26 scheme and the C12 SBR mechanisms based on reaction, species, and RO_2 counts. CRI v2.2 falls between C12multi-L and the more strongly reduced SBR mechanisms in overall mechanism size. Further details are provided in the new Sect. S3.3.3 and Table S13 of the Supplement, with a corresponding reference added to the SBR results section of the revised manuscript:

> *Sect. 6.1.2, page 35, line 747*
>
> Among the PVOCs considered, C12 provides an opportunity to compare the SBR mechanisms with the CRI v2.2 dodecane scheme (Jenkin et al., 2008) (Table S13; Sect. S3.3.3).

Specific Comments

RC: 1.) Introduction: two categories of systematic mechanism reduction are put forth in the introduction: chemistry informed rule-based reduction versus machine learning/data driven reduction. The agglomeration of all non-ML techniques into one category is a simplification that ignores other possibilities for mathematical chemistry-based reductions. In the field of combustion flame kinetics, rigorous mechanism reduction has been developed much earlier than in the field of atmospheric chemistry. Nagy and Turányi (2009) developed the simulation error minimization connectivity method which generates thousands of candidate reduced mechanisms on the basis of the inspection of the normalized Jacobian and produces a database, in which the errors of simulations using the reduced mechanisms are recorded. Based on this database, an optimal

reduced mechanism is selected. The directed relation graph (DRG) method uses a directed graph to map the coupling of specues and consequently find unimportant species for removal based on an acceptable threshold. Sander et al. (2019) have used the graph-oriented reduction method with error propagation by Niemeyer et al. (2010) to reduce the Mainz Organic Mechanism (MOM) and compared the reduced and the full mechanism in a global simulation. For completeness, MOM that is included in the CAABA/ MECCA box model (Sander etal., 2019) should be mentioned as another representative example of detailed VOC mechanisms.

AR: Thank you for the careful reading and detailed explanation. The original two-category description indeed oversimplified the range of systematic mechanism-reduction approaches. This part of the Introduction has therefore been revised to provide a broader overview of chemistry-informed, quantitative kinetic and graph-based, and data-driven approaches, with MOM added as another representative detailed VOC mechanism:

Sect. 1, page 2, line 36

Representative examples include the Master Chemical Mechanism (MCM, Jenkin et al., 1997; Saunders et al., 2003), the Mainz Organic Mechanism (MOM, Lelieveld et al., 2016; Sander et al., 2019), and automated explicit mechanism generators, ...

Sect. 1, page 3, line 54

Systematic ~~reduction approaches for~~ approaches to reducing chemically detailed mechanisms ~~broadly fall into two categories. The first category consists of~~ draw on chemistry-informed ~~, rule-based reduction approaches, which modify detailed mechanisms while retaining mechanistic interpretability, typically by replacing detailed reaction pathways with aggregated or surrogate representations. These methods naturally~~ simplification, quantitative kinetic and graph-based analyses, and data-driven methods. Chemistry-informed approaches evolved from expert-driven manual simplifications ~~. Early applications to VOC mechanisms, including SAPRC (Carter, 2000), MCM (Whitehouse et al., 2004a,b), and GECKO-A (Szopa et al., 2005), incorporate reaction and species grouping or substitution, apply~~ and retain mechanistic interpretability by grouping or substituting species and reactions, applying chemical operators, and ~~use cutoff thresholds for~~ truncating low-yield pathways, particularly ~~for~~ in radical chemistry. ~~Such chemistry-informed~~ These reductions can be ~~implemented during mechanism generation or as a post-processing step~~ incorporated into mechanism construction or generation, as in SAPRC (Carter, 2000) and GECKO-A (Szopa et al., 2005), or applied to an existing mechanism ~~. Other compound-by-compound reductions include the~~ after generation. The Common Representative Intermediates (CRI) scheme ~~, which reduces MCM by grouping intermediate species to preserve key metrics such as~~ similarly derives reduced mechanisms from the MCM on a compound-by-compound basis using representative intermediates selected to preserve O_3 formation potential ~~(Jenkin et al., 2008; Weber et al., 2020). More~~ (Jenkin et al., 2008). Subsequent extensions incorporated selected pathways for highly oxygenated organic molecule formation (Weber et al., 2020).

Quantitative kinetic and graph-based analyses, developed extensively for combustion and pyrolysis mechanisms, identify species and reactions that can be removed or approximated while preserving selected chemical behavior (e.g., Turányi, 1990; Nagy and Turányi, 2009; Niemeyer et al., 2010). In atmospheric applications, the importance of individual species and reactions can similarly be quantified using metrics based on reaction rates, sensitivity and timescale analyses, simulation errors, and kinetically weighted graph dependencies, either individually or in combination. For example, complementary analyses and reduction strategies were combined to reduce the MCM (Whitehouse et al., 2004a,b) and a detailed α-pinene oxidation mechanism (Xia et al., 2009). Sander et al. (2019) applied the directed

relation graph with error propagation (DRGEP) across conditions sampled from a global atmospheric simulation to generate multi-target skeletal versions of MOM. Building on these methodological developments, more recent automated post-generation reduction tools ~~extend this category, operating directly on~~ apply chemistry-informed, quantitative, and graph-based approaches directly to chemically detailed mechanisms~~to systematically reduce reactions and species while preserving targeted chemical behavior~~. The GENerator of reduced Organic Aerosol mechanisms (GENOA) ~~adopts a graph-aware and systematic framework, with evaluation based on simulation feedback on~~ combines chemistry-informed reduction strategies with simulation-based error control, focusing on SOA-relevant chemistry (Wang et al., 2022, 2023), while the Automated ~~Model Reduction for Oxidative Reactions in the Environment~~ MOdel REduction (AMORE) applies graph-theoretical methods to reduce mechanisms with an emphasis on gas-phase chemistry (Wiser et al., 2023, 2025).

~~The second category includes machine-learning~~ Concurrently, data-driven approaches, including machine learning (ML)~~and data-driven approaches, which~~ , provide alternative strategies that replace detailed mechanisms with surrogate representations or emulators trained or optimized for specific applications. For SOA-focused applications, the VBS-GECKO parameterization derives a volatility-basis-set representation with optimized stoichiometries from GECKO-A simulations (Lannuque et al., 2018, 2020). Examples of ML approaches include black-box emulators that ~~reproduce simulation outputs from highly detailed mechanisms using neural networks~~ use neural networks or random forests to reproduce outputs from detailed-mechanism simulations (e.g., Schreck et al., 2022; Mouchel-Vallon and Hodzic, 2023; Kelp et al., 2022), equation-discovery methods such as Sparse Identification of Nonlinear Dynamics (SINDy), which infer reduced ODE systems directly from data (e.g., Yang et al., 2024), and representation-learning techniques such as interpretable autoencoders, which reduce mechanism complexity while reproducing ~~their~~ the time-dependent evolution of the chemical system (e.g., Grassi et al., 2022). While promising, both equation-discovery and representation-learning approaches remain at an early stage of development and have not yet been demonstrated at the level of complexity required for comprehensive VOC oxidation mechanisms.

RC: 2.) Section 2.1: before introducing the highlights of GENOA v3, authors should put forth what the criteria are for measuring the performance of GENOA in this study, in terms of scalability, simulation errors in tSOA, simulation errors in O3 production, mass balance, atom balance of C, O and N, etc.

AR: Thank you for the suggestion. We have clarified the performance criteria at the end of Sect. 2.1 and expanded their description in Sect. 4.5. The specific metrics used in the TBR and SBR analyses are introduced in Sects. 5 and 6, while the SBR acceptance criteria are defined in Sect. 4.4.2 and Appendices C2–C4. As discussed in a previous response, elemental balance is not explicitly enforced during reduction, and related discussion has been added to Sect. 7.2.2. The revised Sect. 2.1 now provides the following overview:

Sect. 2.1, page 4, line 111

The present work details the GENOA v3.0 framework and ~~demonstrates its capabilities for next-generation atmospheric chemistry modeling~~evaluates its performance in terms of mechanism-size reduction, computational cost, and accuracy relative to reference mechanisms for SOA formation and gas-phase chemistry. Originally, GENOA v1 and v2 were collectively referred to as the ...

A summary of the evaluation metrics has also been added to Sect. 4.5:

Sect. 4.5, page 22, line 516

The performance of TBR and SBR ~~performance~~ is presented and discussed in Sects. 5 and 6, respectively. The evaluation ~~focuses on characterizing~~ considers mechanism-size reduction and computational performance, with accuracy assessed using key metrics for SOA and gas-phase chemistry. Mechanism-size reduction is quantified using reaction, species, and condensable-species counts, while computational performance, evaluated primarily for SBR, is assessed using reduction cost and box-model simulation runtime. SOA mass and bulk composition are evaluated using tSOA, elemental ratios, and particle-phase functional-group distributions, while gas-phase behavior is evaluated using bulk reactivities and concentrations of key oxidants and radicals. For SBR, the analysis is extended to temporal variability and dominant reaction pathways. The evaluation focuses on reduction behavior and its relation to chemically interpretable differences among mechanisms~~. Application-ready~~ , while application-ready mechanism development will be the focus of future work.

RC: 3.) Figure 2: Does the gray downwards arrow mean that output of GENOA v1 and v2 can be used as input for GENOA v3?

AR: Thank you for pointing out this ambiguity. The gray downward arrow was intended only to indicate the evolution of GENOA from versions 1 and 2 to version 3, rather than a transfer of mechanism outputs between versions. The arrow has been removed from Fig. 2 (Sect. 3.3.1, page 12) for clarity.

RC: 4.) Section 3.2.2, branch concentration removal: While it makes sense to remove minor branches, a description is missing how the remaining relative branching ratios are corrected after removing the unimportant branch.

AR: Thank you for raising this point. After a branch is removed, the relative branching ratios are recalculated from the retained reactions and renormalized to sum to unity. In the current implementation, this recalculation updates the mechanism-dependent metric used in subsequent TBR steps but does not modify the kinetic or stoichiometric parameters of the retained reactions or redistribute the removed flux. Redistributing this flux while preserving elemental balance could be considered in future development, particularly where cumulative elemental loss is important (discussed in Sect. 7.2.2). We have clarified the current treatment in Sect. 3.2.3:

Sect. 3.2.3, page 11, line 272

~~As a result, the TBR~~ Each time a mechanism-dependent strategy is applied, its reduction metric is recalculated from the current mechanism. These recalculations do not modify the kinetic or stoichiometric parameters of retained reactions or redistribute removed flux. Reaction branching ratios are recalculated over retained reactions with the same reactants and therefore renormalized to sum to unity. The selected strategies are applied iteratively ~~in a predefined order. Strategies that no longer identify reduction candidates are skipped in subsequent passes, and the procedure continues~~ until no further reduction candidates are identified.

RC: 5.) Section 3.3, introduction of SBR: It is stated that the candidate (for the same portion of the mechanism) is selected that achieves the best balance between reduction efficiency and accuracy. How is it ensured that the selected candidate is the best candidate for reduction overall? In other words, how can we be sure that the selected candidate and its subsequent further reduction preclude that a less favorable candidate turn out as better reduction option a few stages downstream? What are "user-defined stopping conditions" and where are they specified?

AR: Thank you for raising these points. As discussed above, SBR is path dependent and does not guarantee a

globally optimal reduction sequence. We agree that "best" and "optimal" could therefore be misleading. We have replaced these terms with "highest-scoring" or "locally optimal", as appropriate, and clarified in Sects. 3.3 and 7.2.3 that candidate selection is local to the current SBR step. Sect. 3.3 has been revised as follows:

Sect. 3.3, page 11, line 288

~~The candidate achieving the best balance between reduction efficiency and accuracy~~ Among the acceptable candidates generated in the current step, the candidate with the highest reduction score is selected and applied~~to the mechanism. This step is repeated iteratively until no further acceptable candidates are available, or until other user-defined stopping conditions are reached~~. As subsequent steps proceed from the selected candidate, this local, path-dependent selection does not guarantee a globally optimal reduction sequence but makes the progressive reduction of large mechanisms computationally tractable.

We have also clarified the cycle stopping/continuation criteria. The original phrase "user-defined stopping conditions" referred to stage settings that determine whether a cycle is repeated or the procedure advances, including an optional minimum number of validated steps required for cycle repetition. These settings and their use in Stages A1 and B1 are detailed in Sect. S3.1. To keep the SBR introduction concise, this phrase was removed from Sect. 3.3, while Appendix B1 (B2 in the preprint) was revised to clarify the definitions of step, cycle, and stage and to specify the cycle stopping/continuation criteria:

Appendix B2, page 54, line 1136

To organize the iterative SBR procedure, GENOA v3 defines three hierarchical levels: step, cycle, and stage. These terms describe how the SBR workflow is applied across species groups and controlled by user-defined configurations.

- Step: A step involves identifying, evaluating, and selecting reduction candidates for a single species group. This is the smallest operational unit in SBR. If an acceptable candidate is found, it is applied to the mechanism and the step is considered validated. ~~Since the change may update the current group, a new step is applied to the same group until it~~ Reduction continues until the current species group becomes empty or no further acceptable ~~reduction is found. The process then continues~~ candidate is found, after which the procedure advances to the next species group.
- Cycle: A cycle is a complete pass through all species groups, ~~performing~~ with one or more steps ~~per groupas needed.~~ performed for each group. Each cycle uses one set of settings specified for the current stage. At the end of ~~each cycle, an optional evaluation process may be used to assess~~ a cycle, the reduced mechanism ~~, with results guiding~~ may optionally be evaluated to guide early termination or ~~prompting adjustments to configuration settings~~adjustments to subsequent settings, as in the pre-testing process used in GENOA v1 and v2.
- Stage: A stage ~~corresponds to a specific configuration defined by the user, including the choice of~~ defines a user-specified SBR configuration, including settings such as reduction strategies, ~~threshold values, and evaluation tolerances . Multiple cycles may be executed within a stage. If a cycle results in no validated steps or too few to meet a continuation criterion, the process~~ training scenarios, and error tolerances (see examples in Sect. 4.4.2). These settings may remain fixed or change in a predefined order across cycles. Once all configured settings have been applied, the

procedure advances to the next stage.

By default, a cycle is repeated with the same settings if it produces at least one validated step. If no steps are validated, the procedure advances to the next ~~stage. SBR ends once~~ settings or, if none remain, to the next stage. Alternatively, users may require a higher minimum number of validated steps before repeating a cycle, as applied in Stages A1 and B1 in this study (Sect. S3.1). This option avoids computationally expensive repetition when another cycle is expected to yield little additional reduction relative to its computational cost, particularly during early stages. SBR ends after all stages have been completed.

RC: 6.) Section 3.3.3, candidate evaluation criteria: elaborate more on the error tolerances. It seems odd to let the user define the error tolerances. Assuming that the user-defined error tolerances are carefully documented in each reduction step, this nevertheless imposes knowledge about prescription of error tolerances on the users. Please describe the objective reasoning for the specific error tolerances (also see my suggestion in the general comment).

AR: We agree that requiring users to define error tolerances may seem unintuitive, since users are ultimately concerned with the realized errors of the reduced mechanism rather than the internal candidate-acceptance bounds. However, SBR requires such bounds to determine whether each candidate can be accepted during iterative reduction. Users do not prescribe the realized error at each reduction step; instead, the three tolerances are specified as stage settings to constrain candidate acceptance as the reduction proceeds.

There is no single objective set of tolerance values because acceptable errors depend on the intended application, evaluation targets, and scenarios. The low-, medium-, and high-tolerance levels used in this study were selected empirically, based on preliminary tests and experience with previous GENOA reductions, to target progressively increasing mean errors of approximately 5, 10, and 15 % and thereby span different size–accuracy trade-offs, rather than to represent universally optimal values. In Sect. 3.3.3, the definition of error tolerances has been rephrased to clarify that they represent acceptance bounds:

Sect. 3.3.3, page 14, line 346

Error tolerances: Thresholds ~~applied to reduction errors that determine whether a candidate is acceptable~~ that define the acceptable bounds for reduction errors. A candidate is considered valid only if all ~~predefined error tolerancesare satisfied~~reduction errors remain within the predefined tolerances.

To provide practical guidance for selecting these settings and to address the related comments on error tolerances, as noted in our previous responses, we added a sensitivity analysis in Sect. S3.1.4 in which the three final tolerance sets were applied independently to U8. The results show that progressively relaxed tolerances allow greater reduction at the expense of larger realized errors, although they do not guarantee the smallest mechanism at each tolerance level because SBR is path dependent. Based on these results, we recommend progressive relaxation as a practical starting strategy, allowing users to evaluate mechanisms obtained under increasingly permissive bounds and select one based on its realized errors. Example configurations are also provided in the GENOA v3 test cases.

We also revised Sect. 4.4.2 to explain how the error tolerances are applied during the staged reduction and to reference this new sensitivity analysis:

Sect. 4.4.2, page 20, line 489

Each stage includes multiple reduction cycles (Appendix B1), with stage-specific settings defined as lists and applied sequentially across cycles (e.g., different error tolerances). Settings may remain fixed or vary within a stage. Here, error tolerances are fixed in the early stages, where repeated cycles already provide substantial reduction. They are progressively relaxed within later stages to track the reduction under increasing allowable errors and obtain mechanisms spanning different size–accuracy trade-offs. These staged tolerance sequences were selected empirically. A complementary U8 test is provided in the Supplement (Sect. S3.1.4), comparing progressive relaxation with direct application of the final low-, medium-, and high-tolerance sets from U8A1 and U8B1 (Table S9).

RC: 7.) Section 4.1, test mechanism selection: Include a reference to Table 6 already here. Consider including the total number of RO2 species in Table 6.

AR: Thank you for these suggestions. We have added a reference to Table 6 in Sect. 4.1:

Sect. 4.1, page 15, line 382

The mechanisms obtained after the preparation step serve as references, and the corresponding TBR mechanisms derived from them are used as starting-point mechanisms for ~~reduction, as outlined~~ SBR. The SBR configurations are described in Sect. 4.4, and the mechanism sizes and reduction accuracy are summarized in Table 6.

Adding another column to the already information-dense Table 6 would reduce its readability while reporting only the final number of RO_2 species in each SBR mechanism. Instead, we added a more comprehensive analysis in the new Sect. S3.3.4 of the Supplement, which reports changes in the numbers of RO_2 species and major VOC and RO_2 reaction classes throughout SBR. References to this analysis were added at relevant points in the discussion:

Sect. 6.1, page 34, line 703

To complement the overall size analysis, Sect. S3.3.4 of the Supplement provides RO_2 species counts and the numbers and relative distributions of selected reaction classes across SBR stages (Tables S14–S15), showing variations with both the PVOC and evaluation targets.

Sect. 6.2.2, page 39, line 829

Complementary structural analysis shows that the final multi-target mechanisms retain more RO_2 species and reactions than the corresponding SOA-focused mechanisms, consistent with their broader representation of HO_x-, NO_x-, and oxidant-coupled chemistry (Table S15; Sect. S3.3.4).

Sect. 6.4, page 42, line 885

Consistent with these pathway comparisons, the LIM multi-target mechanisms retain larger O_3- and NO_3-initiated fractions and more RO_2 species and reactions than the corresponding U8 and C12 mechanisms, indicating broader representation of initiation and peroxy-radical chemistry (Sect. S3.3.4).

Sect. 6.4.3, page 46, line 1051

These results show that SBR selectively preserves chemical pathways according to the evaluation

targets, with SOA-focused reductions favoring compact representations and multi-target reductions retaining broader chemistry at increased complexity, consistent with the structural changes quantified in Sect. S3.3.4.

For completeness, the corresponding reaction-class analysis for TBR is provided in Sect. S2.3 and referenced in the related TBR discussion.

RC: 8.) P16, L 410: Introduce briefly which SOA formation algorithm is applied in the box model simulation. Is the initial SOA seed mass non-volatile? Is the partitioning reversible?

AR: As suggested, we have added information on the SOA formation treatment:

Sect. 4.2, page 16, line 409

SOA formation is simulated using the reversible dynamic gas–particle partitioning scheme of GECKO-A (Aumont et al., 2026), based on Fuchs–Sutugin mass transfer and Raoult's-law absorptive partitioning, with default accommodation and particle-phase activity coefficients of unity. A high initial ~~SOA~~ non-volatile organic aerosol seed mass of 159 μg m^{-3} is prescribed (implemented as 2.23×10^{11} molecules cm^{-3} assuming a molar mass of 427 g mol^{-1}) to ensure efficient gas–particle partitioning.

We have also updated the relevant GECKO-A references throughout the manuscript to cite the newly available model description for the versions of the mechanism generator and box model used in this study.

RC: 9.) Figure 3: why is there no 10% limit shown in plot (b)?

AR: The 10 % reference line is not shown in Fig. 3(b) (Sect. 4.3, page 17) because all OH errors are below approximately 5.5 %, and the line therefore falls outside the displayed y-axis range. We have retained the current axis range to show the variation among the OH errors more clearly.

RC: 10.) Table 5: Explain in the text, why the error tolerances increase within one stage and throughout the stages. P22, L 501: It is mentioned here that error tolerances are relaxed after stage A2 without giving the rationale behind this. P22, L 511: Did you test how sensitive is the reduced mechanism to mean reduction errors larger than 10%, for example in the prediction of tSOA, and how does the choice of larger reduction errors affect the size reduction?

AR: As discussed in our responses to the general comment and specific comment 6, the staged tolerances were selected empirically to span representative size–accuracy trade-offs. In this study, the tolerances were fixed in the early stages, where repeated cycles already provided substantial reduction, and progressively relaxed within Stages A2 and B2 to generate mechanisms at different allowable error levels. Stage A3 does not continue this sequence; instead, its branches originate from different A2 mechanisms and use tolerances comparable to or slightly relaxed from those of the corresponding A2 mechanisms. We also added a U8 sensitivity test in Sect. S3.1.4, including a mean-error tolerance above 10 %. The results show that relaxing the tolerances generally permits greater size reduction at the expense of larger realized errors. They also show that the resulting mechanism depends on the tolerance sequence because SBR is path dependent.

RC: 11.) P24, L 550: What are the criteria to consider a tSOA error in TBR as acceptable? That should be defined from the outset and may also be different for different PVOC. Related to this (P26, L 583), it is shown that tSOA in T10 is systematically giving 10-20% lower SOA mass than the Reference. This potentially leads to a systematic underestimation of observed SOA in simulation with a box model using this reaction mechanism. Please comment on that issue and how it can be circumvented.

AR: This criterion is described in Sect. 4.3 (line 445) as part of the TBR configuration. An error level of approximately 10 % was used for all evaluation metrics, including tSOA. This level was not defined as a formal acceptability criterion but as a simple empirical guide for screening the overall size–accuracy trade-off. As the reviewer notes, an appropriate error level depends on the PVOC and intended application and should therefore be evaluated for each use case. We have clarified this point in Sect. 4.3:

> *Sect. 4.3, page 18, line 445*
>
> Based on these analyses, suitable threshold values were selected for each strategy ~~that achieve~~ to provide relatively large size ~~reduction while maintaining acceptable accuracy across mechanisms and evaluation targets (~~reductions while keeping errors near 10 %~~), which served as an empirical guide for screening the overall size–accuracy trade-off.

We also agree that the predominantly negative T10 bias indicates over-reduction for some PVOCs. This bias is relative to the reference mechanism and would lead to additional underestimation relative to observations only where the reference mechanism accurately reproduces the observed SOA. As discussed in Sects. 5.2.1, 5.3.2, and 7.1, the current uniform T10 configuration is not suitable for all mechanisms. Such underestimation can be limited by using the more conservative T6 configuration or by adjusting the applied strategies and thresholds for the target PVOC and application.

RC: 12.) P31, L679-686: A possible conclusion at the end of the TBR performance evaluation could be that T6 can be used for multitarget reduction (gas-phase and SOA related) and T10 only for SOA-related reductions. This also could lead to the conclusion that it is better to streamline reductions with TBR by just allowing two different options, T6 or T10, with a given set of pre-defined reduction parameters for each.

AR: Thank you for this insightful suggestion. We have revised the conclusion of the TBR performance evaluation accordingly:

> *Sect. 5.3.2, page 31, line 679*
>
> Overall, gas-phase reactivity metrics (Figs. 10 and~~S19~~ S20) show consistent patterns and are aligned with the SOA-related results, indicating that TBR reduction preserves dominant chemical pathways, with performance depending on mechanism size and chemical complexity. ~~Under~~ The small deviations under T6 ~~, deviations in~~ support its use when both gas-phase and SOA-related ~~metrics remain small, whereas under~~ chemistry are to be preserved. T10 ~~larger deviations emerge, reflecting increased sensitivity to over-reduction and to the cumulative removal of distributed pathways~~provides greater size reduction but introduces larger deviations, making it an efficient initial reduction option for SOA-focused applications when its thresholds are tailored to the target mechanism. The metric evaluation also suggests that the current uniform ~~, environmentally sensitive strategy~~ strategies and threshold choices can be further refined for small and chemically complex mechanisms. In such cases, applying additional mechanism-specific constraints targeted to selected gas-phase or SOA-related metrics, or adjusting simulation settings, would improve related preservation.

RC: 13.) Section 7.2.1, scenario dependence: Not spanning scenarios that occur under certain atmospheric conditions put the mechanism reduction at risk, although the selected 11 scenarios cover most typical environments in the lower atmosphere. An additional scenario should include light-dependent emissions of isoprene, as mixing with isoprene may potentially suppress SOA and this interference is important everywhere in the continental atmosphere. Moreover, a scenario representing conditions of the upper troposphere in the tropics would assist in the reduction for use in global chemistry-transport models.

AR: Thank you for these suggestions. Both scenarios would broaden the evaluation beyond the controlled single-precursor, lower-tropospheric conditions considered in this study and are relevant for future application-oriented reductions. We have clarified these limitations and the need for appropriately tailored scenario sets in the discussion:

> *Sect. 7.2.2, page 47, line 991*
>
> The 11 scenarios target lower-tropospheric VOC oxidation (Sect. 4.2) but do not represent potentially important precursor interactions, including nonlinear SOA responses in VOC mixtures (McFiggans et al., 2019). Additionally, reductions targeting distinct atmospheric regimes (e.g., the tropical upper troposphere) or intended for use in global CTMs would require correspondingly tailored scenario sets (Sander et al., 2019).

RC: 14.) Section 7.4, integration with machine learning: Although GENOA retains mechanistic interpretability, neither ML-based approaches nor GENOA allow tracking of reduction decisions to find out why a certain reduction candidate was preferred over others. With GENOA there is at least theoretically a chance to trace a reduction decision, but I guess it is terribly complicated to trace it back after several reduction stages. Please comment.

AR: We agree that mechanistic interpretability and decision traceability are distinct. GENOA preserves all intermediate accepted mechanisms from the starting-point to the final reduced mechanism, together with the reduction scores and simulation errors used for candidate selection (e.g., the U8 reduction trajectory in Sect. S3.1.2 of the Supplement). These intermediate mechanisms may serve as starting points for further reduction or as final mechanisms representing different size–accuracy trade-offs (e.g., U8B2a–U8B2d). The selected pathway and individual decisions can therefore be reconstructed, although tracing them across many steps and stages may be complex and time-consuming.

The broader limitation is related to SBR path dependence: selecting a different candidate at any step may produce a different reduction pathway and final mechanism, while systematically exploring alternative branches would substantially increase the computational cost. This issue has already been clarified in revised Sect. 7.2.3 and addressed in our previous responses.

References

Aumont, B., Valorso, R., Rickard, A., Jenkin, M., Lee-Taylor, J., Orlando, J., Madronich, S., Vereecken, L., and Camredon, M.: GECKO-A v1.0: Exploring VOC Oxidation Trajectories Through Comparison with the Master Chemical Mechanism, EGUsphere, 2026, 1–47, , 2026.

Carter, W. P. L.: Documentation of the SAPRC-99 Chemical Mechanism for VOC Reactivity Assessment, , 2000.

Grassi, T., Nauman, F., Ramsey, J., Bovino, S., Picogna, G., and Ercolano, B.: Reducing the complexity of chemical networks via interpretable autoencoders, Astronomy & Astrophysics, 668, A139, , 2022.

Jenkin, M., Watson, L., Utembe, S., and Shallcross, D.: A Common Representative Intermediates (CRI) mechanism for VOC degradation. Part 1: Gas phase mechanism development, Atmos. Environ., 42, 7185–7195, , 2008.

Jenkin, M. E., Saunders, S. M., and Pilling, M. J.: The tropospheric degradation of volatile organic compounds: a protocol for mechanism development, Atmos. Environ., 31, 81–104, , 1997.

Kelp, M. M., Jacob, D. J., Lin, H., and Sulprizio, M. P.: An online-learned neural network chemical solver for stable long-term global simulations of atmospheric chemistry, J. Adv. Model. Earth Syst., 14, e2021MS002 926, , 2022.

Lannuque, V., Camredon, M., Couvidat, F., Hodzic, A., Valorso, R., Madronich, S., Bessagnet, B., and Aumont, B.: Exploration of the influence of environmental conditions on secondary organic aerosol formation and organic species properties using explicit simulations: development of the VBS-GECKO parameterization, Atmos. Chem. Phys., 18, 13 411–13 428, , 2018.

Lannuque, V., Couvidat, F., Camredon, M., Aumont, B., and Bessagnet, B.: Modeling organic aerosol over Europe in summer conditions with the VBS-GECKO parameterization: sensitivity to secondary organic compound properties and IVOC (intermediate-volatility organic compound) emissions, Atmos. Chem. Phys., 20, 4905–4931, , 2020.

Lelieveld, J., Gromov, S., Pozzer, A., and Taraborrelli, D.: Global tropospheric hydroxyl distribution, budget and reactivity, Atmos. Chem. Phys., 16, 12 477–12 493, , 2016.

McFiggans, G., Mentel, T. F., Wildt, J., Pullinen, I., Kang, S., Kleist, E., Schmitt, S., Springer, M., Tillmann, R., Wu, C., et al.: Secondary organic aerosol reduced by mixture of atmospheric vapours, Nature, 565, 587–593, , 2019.

Mouchel-Vallon, C. and Hodzic, A.: Toward emulating an explicit organic chemistry mechanism with random forest models, J. Geophys. Res.-Atmos., 128, e2022JD038 227, , 2023.

Nagy, T. and Turányi, T.: Reduction of very large reaction mechanisms using methods based on simulation error minimization, Combust. Flame, 156, 417–428, , 2009.

Niemeyer, K. E., Sung, C.-J., and Raju, M. P.: Skeletal mechanism generation for surrogate fuels using directed relation graph with error propagation and sensitivity analysis, Combust. Flame, 157, 1760–1770, , 2010.

Sander, R., Baumgaertner, A., Cabrera-Perez, D., Frank, F., Gromov, S., Grooß, J.-U., Harder, H., Huijnen, V., Jöckel, P., Karydis, V. A., Niemeyer, K. E., Pozzer, A., Riede, H., Schultz, M. G., Taraborrelli, D., and Tauer, S.: The community atmospheric chemistry box model CAABA/MECCA-4.0, Geosci. Model Dev., 12, 1365–1385, , 2019.

Saunders, S. M., Jenkin, M. E., Derwent, R., and Pilling, M.: Protocol for the development of the Master Chemical Mechanism, MCM v3 (Part A): tropospheric degradation of non-aromatic volatile organic compounds, Atmos. Chem. Phys., 3, 161–180, , 2003.

Schreck, J. S., Becker, C., Gagne, D. J., Lawrence, K., Wang, S., Mouchel-Vallon, C., Choi, J., and Hodzic, A.: Neural Network Emulation of the Formation of Organic Aerosols Based on the Explicit GECKO-A Chemistry Model, J. Adv. Model. Earth Syst., 14, e2021MS002 974, , 2022.

Szopa, S., Aumont, B., and Madronich, S.: Assessment of the reduction methods used to develop chemical schemes: building of a new chemical scheme for VOC oxidation suited to three-dimensional multiscale HO_x-NO_x-VOC chemistry simulations, Atmos. Chem. Phys., 5, 2519–2538, , 2005.

Turányi, T.: Reduction of large reaction mechanisms, New J. Chem., 14, 795–803, `https://api.semanticscholar.org/CorpusID:99253656`, 1990.

Wang, Z., Couvidat, F., and Sartelet, K.: GENerator of reduced Organic Aerosol mechanism (GENOA v1.0): an automatic generation tool of semi-explicit mechanisms, Geosci. Model Dev., 15, 8957–8982, , 2022.

Wang, Z., Couvidat, F., and Sartelet, K.: Implementation of a parallel reduction algorithm in the GENerator of reduced Organic Aerosol mechanisms (GENOA v2.0): Application to multiple monoterpene aerosol precursors, J. Aerosol Sci., 174, 106 248, , 2023.

Wang, Z., Couvidat, F., and Sartelet, K.: Response of biogenic secondary organic aerosol formation to anthropogenic NOx emission mitigation, Sci. Total Environ., 927, 172 142, , 2024.

Weber, J., Archer-Nicholls, S., Griffiths, P., Berndt, T., Jenkin, M., Gordon, H., Knote, C., and Archibald, A. T.: CRI-HOM: A novel chemical mechanism for simulating highly oxygenated organic molecules (HOMs) in global chemistry–aerosol–climate models, Atmos. Chem. Phys., 20, 10 889–10 910, , 2020.

Whitehouse, L. E., Tomlin, A. S., and Pilling, M. J.: Systematic reduction of complex tropospheric chemical mechanisms, Part I: sensitivity and time-scale analyses, Atmos. Chem. Phys., 4, 2025–2056, , 2004a.

Whitehouse, L. E., Tomlin, A. S., and Pilling, M. J.: Systematic reduction of complex tropospheric chemical mechanisms, Part II: Lumping using a time-scale based approach, Atmos. Chem. Phys., 4, 2057–2081, , 2004b.

Wiser, F., Place, B. K., Sen, S., Pye, H. O. T., Yang, B., Westervelt, D. M., Henze, D. K., Fiore, A. M., and McNeill, V. F.: AMORE-Isoprene v1.0: a new reduced mechanism for gas-phase isoprene oxidation, Geosci. Model Dev., 16, 1801–1821, , 2023.

Wiser, F., Sen, S., Wang, Z., Lee-Taylor, J., Barsanti, K. C., Orlando, J., Westervelt, D. M., Henze, D. K., Fiore, A. M., Berman, A., Carter, R., and McNeill, V. F.: A graph theory-based algorithm for the reduction of atmospheric chemical mechanisms, PNAS Nexus, 4, pgaf273, , 2025.

Xia, A. G., Michelangeli, D. V., and Makar, P. A.: Mechanism reduction for the formation of secondary organic aerosol for integration into a 3-dimensional regional air quality model: α-pinene oxidation system, Atmos. Chem. Phys., 9, 4341–4362, , 2009.

Yang, X., Guo, L., Zheng, Z., Riemer, N., and Tessum, C. W.: Atmospheric chemistry surrogate modeling with sparse identification of nonlinear dynamics, Journal of Geophysical Research: Machine Learning and Computation, 1, e2024JH000 132, , 2024.