https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

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

Zhizhao Wang[1,2], William P. L. Carter[1], Julia Lee-Taylor[2], John Orlando[2], Qing Ye[2], Richard Valorso[3], Marie Camredon[3], Bernard Aumont[3], and Kelley Barsanti[2]

[1]College of Engineering Center for Environmental Research and Technology (CE-CERT), University of California, Riverside, CA 92521, USA
[2]Atmospheric Chemistry Observations & Modeling Lab (ACOM), NSF National Center for Atmospheric Research, Boulder, CO 80301, USA
[3]Univ Paris Est Creteil and Université Paris Cité, CNRS, LISA, F-94010 Créteil, France

Correspondence: Zhizhao Wang (zhizhaow@ucr.edu) and Kelley Barsanti (barsanti@ucar.edu)

Abstract.

Comprehensive atmospheric chemical mechanisms for volatile organic compound (VOC) oxidation contain thousands to millions of reactions and species, presenting major computational challenges for large-scale or long-term simulations. As mechanism complexity continues to increase, reduction strategies are required to enable their use in atmospheric modeling while preserving accuracy.

This paper presents the GENerator of Optimized Atmospheric chemical mechanisms (GENOA v3), a major advancement over earlier GENOA versions that enables scalable reduction of highly detailed mechanisms containing up to millions of reactions and species. GENOA v3 combines fast, strategy-driven threshold-based reduction (TBR) with simulation-based reduction (SBR) that explicitly controls accuracy. The framework is modular, graph-aware, and user-configurable, resulting in compact and chemically interpretable reduced mechanisms.

Applications to GECKO-A mechanisms for diverse VOC precursors across a range of scenarios show that TBR achieves mechanism size reductions of 20–90 % while preserving reasonable accuracy for metrics related to secondary organic aerosol (SOA) formation and gas-phase chemistry, with performance systematically dependent on precursor structure and chemical complexity across mechanisms. SBR achieves further reductions in mechanism size by several orders of magnitude; when trained with 15 % mean error constraints, SBR produces schemes within 0.02 % of the original size for preservation of SOA mass and 0.05 % when also considering gas-phase reactivity (e.g., OH, O_3, and NO_3).

These results demonstrate for the first time that GENOA v3 can reduce highly detailed chemical mechanisms while jointly preserving SOA mass and gas-phase reactivity, achieving substantial size reductions with reasonable accuracy across a wide range of scenarios. Continued application of GENOA v3 and growth of a user community will potentially support the development of libraries of reduced mechanisms and optimized reduction strategy sets tailored to specific modeling applications.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

1 Introduction

Volatile organic compounds (VOCs), emitted from both anthropogenic and biogenic sources, significantly influence atmospheric chemistry through complex oxidation and degradation processes, impacting human health, climate, and ecosystems (e.g., Atkinson and Arey, 2003; Lelieveld et al., 2015; Mouchel-Vallon et al., 2020). Atmospheric processing of these VOCs forms numerous reactive intermediates and stable organic compounds spanning multiple oxidation states, whose levels and reactivity are largely influenced by ambient conditions, including photochemical activity and pollutant regimes (Vereecken and Francisco, 2012; Porter et al., 2021). One critical atmospheric impact of VOC oxidation is the formation of secondary pollutants, including secondary organic aerosol (SOA), through the production of oxidation products spanning a wide range of volatilities, from semi-volatile and non-volatile organic compounds (SVOCs and NVOCs), which can partition to the particle phase under favorable conditions (Hallquist et al., 2009). Moreover, the processing of VOCs can substantially modify atmospheric oxidative capacity by influencing the concentrations and chemical processes of reactive species, including ozone (O_3), the hydroxyl radical (OH), the hydroperoxy radical (HO_2), the nitrate radical (NO_3), and nitrogen oxides (NO_x), thereby altering atmospheric reactivity and air quality (Seinfeld and Pandis, 2016).

Chemically detailed mechanisms have been developed to capture how VOC oxidation pathways and product formation depend on molecular structure and ambient conditions, including substantial variations across daytime and nighttime, urban and rural, and different NO_x regimes (Orlando and Tyndall, 2012; Ehn et al., 2014; Bianchi et al., 2019). Representative examples include the Master Chemical Mechanism (MCM, Jenkin et al., 1997; Saunders et al., 2003), and automated explicit mechanism generators, including the Generator for Explicit Chemistry and Kinetics of Organics in the Atmosphere (GECKO-A, Aumont et al., 2005; Camredon et al., 2007; Mouchel-Vallon et al., 2020) and the SAPRC Mechanism Generation System (MechGen, Carter et al., 2025a, b). The development of these chemically detailed mechanisms and generators relies on structure-activity relationship (SAR) protocols, which are continuously improved to estimate reaction rates and product yields in the absence of comprehensive experimental data (e.g., Kerdouci et al., 2014; Jenkin et al., 2018, 2020; Carter, 2021; Carter et al., 2025a). As new SARs are developed and implemented (e.g., Newland et al., 2022), these chemically detailed mechanisms continue to expand in scope and complexity (e.g., Roldin et al., 2019), increasing our ability to systematically represent the complexity of gas-phase chemistry of atmospheric oxidants and pollutants, including its role in SOA formation, under diverse conditions and for different compound types in air quality simulations.

Although detailed chemical mechanisms greatly enhance our understanding of atmospheric chemistry and may improve modeling, their complexity with potentially thousands or millions of reactions and species poses significant computational challenges, especially in large-scale or long-term atmospheric simulations (Kanakidou et al., 2005; Stockwell et al., 2012; Ervens et al., 2024). In such applications, where computational constraints require simplified, lumped representations of VOC oxidation chemistry, top-down mechanism reduction techniques have become essential. Mechanism reduction aims to simplify detailed chemical schemes while preserving the essential chemical pathways and consistency with the original mechanism, which is critical for predicting key atmospheric constituents and processes, including SOA production and oxidant levels.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

Systematic reduction approaches for chemically detailed mechanisms broadly fall into two categories. The first category consists of 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 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 chemical operators, and use cutoff thresholds for low-yield pathways, particularly for radical chemistry. Such chemistry-informed reductions can be implemented during mechanism generation or as a post-processing step applied to an existing mechanism. Other compound-by-compound reductions include the Common Representative Intermediates (CRI) scheme, which reduces MCM by grouping intermediate species to preserve key metrics such as O_3 formation potential (Jenkin et al., 2008; Weber et al., 2020). More recent automated post-generation reduction tools extend this category, operating directly on 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 SOA-relevant chemistry (Wang et al., 2022, 2023), while the Automated Model Reduction for Oxidative Reactions in the Environment (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 (ML) and data-driven approaches, which 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 (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 time-dependent evolution (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.

In this study, we introduce GENOA v3, the latest and most advanced version of the GENOA algorithm. While retaining the fundamental top-down reduction framework of previous versions, GENOA v3 significantly expands its capabilities to operate on highly detailed chemical mechanisms with up to millions of species and reactions, and to preserve not only SOA mass but also broader gas-phase chemistry indicators and key evaluation metrics. The goal is to provide a flexible and extensible reduction tool that supports both chemistry-informed and data-driven strategies, while allowing user customization and future methodological extensions. The GENOA tool is briefly introduced in Sect. 2, along with the new features of GENOA v3. Detailed reduction descriptions are presented in Sect. 3. Model evaluation is performed using various detailed VOC mechanisms generated by GECKO-A, with setups and results presented in Sects. 4 to 6. The evaluation includes reduction performance across multiple chemical targets, extending beyond SOA mass to include oxidants and other pollutants. Finally, conclusions are drawn in Sect. 8.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

2 Model overview

2.1 The GENOA versions

The GENOA tool comprises a family of reduction algorithms, currently three versions (GENOA v1, v2, and v3), developed to generate condensed chemical mechanisms from detailed atmospheric VOC degradation schemes.

GENOA v1 (Wang et al., 2022) adopted a series-reduction framework, focusing on preserving pathways related to SOA formation. It was successfully applied to the MCM sesquiterpene degradation scheme, reducing the mechanism from 1 625 reactions and 579 species to just 23 reactions and 15 species, with a mean error below 3 % in the simulated total SOA mass (tSOA) relative to the full mechanism.

Building on v1, GENOA v2 (Wang et al., 2023) introduced a parallel-reduction framework, enabling the reduction of mechanisms involving multiple biogenic and anthropogenic precursors for both regional- (Wang et al., 2024) and urban- (Sartelet et al., 2024) scale modeling. This framework explores multiple reduction possibilities for a given subset of the mechanism in parallel, and applies a scoring system to identify the optimal candidate among the alternatives. In one application, a mechanism including 3 001 reactions and 1 227 species for the oxidation of three monoterpenes was reduced to 197 reactions and 110 species, with a mean tSOA error of approximately 3 %. Both versions were integrated with the SSH-aerosol model (Sartelet et al., 2020), enabling treatment of non-ideality and hydrophilic/hydrophobic partitioning during gas-to-particle partitioning.

While GENOA v1 and v2 efficiently reduced mechanisms containing several thousand reactions and species, their scalability was limited. As the mechanism size increased, overall reduction time scaled nonlinearly for both versions, and memory issues became overwhelming when input mechanisms exceeded ten thousand reactions. These limitations rendered GENOA v1 and v2 unsuitable for use with highly detailed mechanisms, such as those generated by GECKO-A, which can contain millions of reactions and species.

To overcome these challenges, GENOA v3 introduces major advances in reduction efficiency and scalability, enabling reductions of highly detailed chemical mechanisms that were previously inaccessible. Table 1 summarizes the key differences among all versions. The present work details the GENOA v3.0 framework and demonstrates its capabilities for next-generation atmospheric chemistry modeling. Originally, GENOA v1 and v2 were collectively referred to as the "GENerator of reduced Organic Aerosol mechanisms". To reflect its expanded capability and broader applicability in GENOA v3, the acronym has been redefined as "GENerator of Optimized Atmospheric chemical mechanisms", highlighting its evolution into a general-purpose framework for mechanism reduction and exploration. The GENOA v3.0 archive is available on Zenodo (Wang, 2026).

2.2 Highlights of GENOA v3

Major developments in this new version include:

- Compatibility with detailed mechanisms: GENOA v3 supports the reduction of highly detailed VOC mechanisms generated by GECKO-A, thereby allowing GENOA to handle much larger mechanisms than previous versions. The interfaces between GENOA and the GECKO-A box model have been updated to incorporate mechanism evaluation based on user-

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

Table 1. An overview of the key differences across GENOA versions.

Feature	v1	v2	v3
Mechanism scale [a]	$\sim 10^3$	$\sim 10^3$	$> 10^6$
Mechanism format [b]	MCM	MCM	GCK, MCM
Multiple precursors [c]	No	Yes	Yes
Parallel reduction [d]	No	Yes	Yes
Dynamic reduction [e]	No	No	Yes
Evaluation target [f]	tSOA	User	User
Box model coupling [g]	SSH	SSH	GCK, SSH
Code extensibility [h]	Low	Low	High

[a] Approximate number of reactions and species supported.
[b] MCM: Master Chemical Mechanism; GCK: GECKO-A mechanism.
[c] Supports reductions of mechanisms from multiple precursors.
[d] Supports parallel identification and evaluation among alternative reduction possibilities. [e] Enables dynamic identification and evaluation of reduction possibilities (Sect. 3.3.3).
[f] Targets preserved during reduction. "User" indicates support for multiple user-defined evaluation targets (e.g., including tSOA and gas-phase tracers). This feature was implemented in v2 and tested in v3.
[g] SSH: SSH-aerosol box model; GCK: GECKO-A box model.
[h] Code is modular and structured to support future development.

customizable metrics, ensuring that VOC degradation via reduced mechanisms preserves chemical indicators (e.g., SOA formation, oxidant levels) across a range of ambient or chamber scenarios.

- Improved reduction workflow and scalability: GENOA v3 introduces an enhanced reduction workflow that improves accuracy, flexibility, and computational scalability for large and complex chemical mechanisms. Its reduction workflow, consisting of two processes that differ in scope, computational cost, and level of detail, is a central focus of this paper and is described in detail in Sect. 3.

- Integrated analysis tools and user accessibility: GENOA v3 includes functionalities for mechanism interpretation (e.g., parsing or modifying MCM and GECKO-A mechanisms), 0-D chemical box simulations (e.g., using the GECKO-A box model or SSH-aerosol), and post-processing analysis (e.g., visualizing chemical pathways and simulation outputs). These tools operate independently of the reduction routine and support both modeling and experimental applications. A static web-based documentation site provides basic information and test cases, assisting users in applying GENOA v3 and its associated tools.

These changes are supported by a complete refactoring of the Python codebase into a modular framework, improving clarity, efficiency, and long-term usability.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

3 Reduction representation

GENOA v3 reduces the complexity of atmospheric chemical mechanisms by identifying and applying reduction candidates, defined as species and their associated reactions targeted to decrease overall mechanism size and complexity.

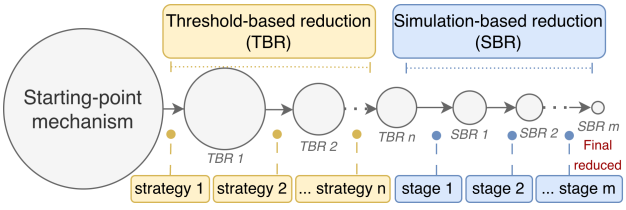

Figure 1. Illustration of the reduction workflow in GENOA v3. TBR n and SBR m represent mechanisms obtained after the nth threshold-based step or the mth simulation-based stage.

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 two main processes:

- 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.

- 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.

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.

The mechanism representation that underpins these processes is described in Sect. 3.1, and detailed workflows for TBR and SBR are presented in Sects. 3.2 and 3.3, respectively. Table 2 provides a concise reference for key terminology used in GENOA v3, with more detailed explanations discussed throughout this section. Appendix B5 provides an illustrative example using a schematic mechanism to support explanation and visualization of key concepts in reduction representation.

3.1 Mechanism representation

Before reduction begins in GENOA v3, the input chemical mechanism is parsed into two main components: a species list and a reaction list. These components are used to build structured, object-based representations of species, reactions, and kinetics that

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

Table 2. Summary of key terminology used in GENOA v3 reduction.

Mechanism and reduction context	
Reference mechanism [a]	Mechanism selected as the baseline for comparisons.
Starting-point mechanism	Mechanism on which the reduction is performed.
Reducible species	Organic species that are eligible for reduction.
Reducible reaction	Reaction containing reducible species that is eligible for reduction.
Reference concentration	Concentration profiles used to derive reduction-relevant information (e.g., lifetime).
Threshold-based reduction (TBR)	GENOA v3 reduction guided by rule-driven thresholds.
Simulation-based reduction (SBR)	GENOA v3 reduction guided by 0-D simulation feedback.
Reduction identification	
Reduction strategy	Protocols defining how species or reactions may be modified. Applied strategies differ between TBR and SBR.
Reduction candidate	Proposed potential modification generated by applying a reduction strategy to the full mechanism (TBR) or a subset (SBR).
Group list (SBR only [b])	Ordered grouping of reducible species used to define subsets of the mechanism.
Base candidate (SBR only)	Candidate generated from one group and one strategy.
Variant candidate (SBR only)	Subset or altered form of the base candidate.
Reduction evaluation	
Evaluation target	Species, species groups, or derived quantities selected to evaluate accuracy.
Evaluation scenario	Scenario selected to test candidate performance.
Error measure	Measure quantifying deviation from the reference mechanism.
Reduction score (SBR only)	Measure describing the size reduction achieved by a candidate.
Error tolerance (SBR only)	Acceptable bound applied to the error measure.

[a] Multiple reference mechanisms are possible; one is used in this work for clarity; [b] concepts used only in SBR.

facilitate the reduction process. Formal definitions of the data structures, along with summaries of key properties for species and reactions, are provided in Appendix A.

3.1.1 Species object

Chemical species in the mechanism are classified as either reducible or non-reducible, depending on their role within the mechanism. Non-reducible species include inorganic compounds and small organic molecules involved in basic reactions. These species are typically the same from any input mechanism, as they participate in fundamental or background chemistry that is not subject to reduction. Reducible species, on the other hand, are the remaining organic compounds formed from the oxidation and degradation of target VOCs. This group varies by input mechanism and includes intermediate and product organic compounds that can be modified or removed during the reduction process.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

Each species is represented as an object, identified by a unique name and associated with several chemical and physical properties. These properties are utilized in the reduction workflow for reducible species, as defined in Appendix A1 and summarized in Table A1, and include molecular weight, molecular formula ($C_xH_yO_zN_w$), SMILES structure, chemical stability (radical or closed-shell), condensability, saturation vapor pressure P_{sat}, and functional group information. Some values are directly obtained from the input files, while others are calculated from their SMILES structures using external tools, such as UManSysProp (Topping et al., 2016); the properties can be updated as needed.

3.1.2 Reaction and kinetic objects

A reaction is defined by its reactants and products, along with their stoichiometric coefficients. Each reaction is also associated with a kinetic object that stores information such as the rate constant and reaction type. The corresponding object structures and kinetic parameters are described in Appendix A2. Basic reactions composed only of non-reducible species are labeled as non-reducible and are not subject to reduction, while other reactions involving reducible species are labeled as reducible and are eligible for reduction.

GENOA v3 reads and parses all reaction types supported by GECKO-A. For unimolecular reactions, the reacting species itself is the reduction target. For most bimolecular reactions, the organic reactant is treated as the targeted reducible species, while the co-reacting species (e.g., OH, O_3, NO_3, NO, HO_2, NO_2) is treated as non-reducible. Organic peroxy radical (RO_2) permutation reactions (RO_2 + RO_2 reactions) are treated following the GECKO-A scheme, using up to nine RO_2 reactivity classes, with each reaction represented by a single targeted RO_2 species reacting with its corresponding class.

To compute species-specific metrics such as lifetimes, branching ratios, and product contributions, reaction rate estimates are required. GENOA v3 evaluates reducible bimolecular reactions as pseudo-unimolecular reactions of the targeted reducible species, allowing their rates to be computed directly. Under this assumption, reaction rates are evaluated using the reference concentration of the co-reacting species, with prescribed environmental parameters applied where relevant (e.g., for temperature- or photolysis-dependent reactions). Reference concentrations are defined as mean values for each species, derived from 0-D simulations of the reference mechanism across a set of selected scenarios. Environmental parameters, including temperature (T), solar zenith angle (SZA), and relative humidity (RH), are specified independently and may be provided as fixed values or as lists. When provided as lists, GENOA v3 evaluates rate ratios across all parameter combinations, and the resulting range is used for reduction. For RO_2 + RO_2 reactions, the concentration of each RO_2 class is defined as the sum of the reference concentrations of its member RO_2 species, allowing these reactions to be treated consistently using the same pseudo-unimolecular approach.

3.1.3 Reaction graph representation

A graph-based representation (e.g., Silva and Halappanavar, 2024; Li et al., 2025) is implemented to capture the interactions between species and reactions, in which nodes represent reducible species, and edges represent reactions connecting reactant nodes to product nodes. The entire mechanism is mapped onto a single unified graph for reduction (e.g., Fig. B1). For each reducible species, the node depth is defined as the shortest path length from any primary VOC (PVOC) in the graph. Each

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

species is further associated with graph-based connectivity sets, which collect all species connected to that species through any reaction pathways, including both upstream precursors and downstream products. Based on these connections, a species branch is defined for each reducible species as the set consisting of the species itself, treated as the branch root, together with all downstream species. These graph-derived properties underpin the graph-based reduction strategies implemented in GENOA v3, with related definitions provided in Appendix A3.

3.2 Threshold-based reduction (TBR)

TBR, similar in function to the pre-reduction steps in GENOA v1 and v2, conducts an initial, mechanism-wide reduction based on user-defined protocols and thresholds. Reduction is driven by rule-based strategies that screen species and reactions for potential removal or other modification according to predefined thresholds. TBR strategies are categorized as either environmentally insensitive, relying solely on molecular or structural properties, or environmentally sensitive, where outcomes depend on conditions such as oxidant levels or photolysis rates. Table B1 provides an overview of the available strategies and their mathematical formulations.

3.2.1 Environmentally insensitive strategies

These strategies are empirical filters targeting reducible components according to intrinsic physical, chemical, or pathway-related properties that do not vary significantly across conditions. They are summarized below, with threshold notation provided:

- NVOC reclassification: SVOCs with P_{sat} below $P_{\text{sat}}^{\text{nvoc}}$ are reclassified as NVOCs. This simplifies gas-particle partitioning by decreasing the number of SVOCs retained in the mechanism.
- VOC reclassification: SVOCs with P_{sat} above $P_{\text{sat}}^{\text{svoc}}$ are reclassified as VOCs. When combined with NVOC reclassification, this allows selective retention of only the most relevant SVOCs for evaluation scenarios.
- NVOC downstream removal: Reactions involving downstream species originating from NVOCs with P_{sat} below $P_{\text{sat}}^{\text{nvoc,dwn}}$ are removed, while the NVOC species itself is retained, reflecting strong partitioning toward the condensed phase and negligible downstream gas-phase chemistry.
- VOC branch removal: Species branches are removed when the root species and all downstream species have P_{sat} above $P_{\text{sat}}^{\text{svoc,brh}}$. Such branches are typically of low relevance for SOA formation, while their importance for gas-phase chemistry is mechanism-dependent.
- Node depth limit: Species with node depth beyond $d^{\max}$ are removed, as species further downstream from the PVOCs are assumed to have a minor influence on dominant chemical pathways. It is similar to the generation number used in GECKO-A, which tracks species by VOC oxidation steps and groups intermediates with different node depths into the same generation.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

– Yield-to-depth ratio removal: Relative effective yields are defined for products formed from reactions sharing the same reactants and reaction type (defaulting to a value of one when no comparison is possible; Appendix B1). Products with relative yield normalized by node depth below $\left(\frac{Y^{\text{rel}}}{d}\right)_{\text{min}}$ are considered removable, accounting for the decreasing importance of downstream species. This cutoff is conceptually similar to the relative-yield–only cutoff commonly applied during mechanism generation (e.g., in GECKO-A), which excludes products with very low relative yields.

– Fast-degradation lifetime jumping: Species with lifetimes shorter than τ^{fast} are considered fast-degrading and bypass explicit representation in the mechanism. This criterion is applied to species whose degradation is dominated by a single pathway, such that their branching ratios are not strongly sensitive to environmental conditions. These species are therefore short-lived across regimes (e.g., alkoxy radicals with lifetimes up to 10^{-4} seconds, Orlando et al. (2003)).

Among the seven environmentally insensitive strategies implemented, four rely on volatility-based thresholds, while the remaining three are based on structural or kinetic properties. Because volatility-based classification is typically applied during mechanism generation, reclassification in GENOA leads to additional reduction only when tighter P_{sat} thresholds are applied to further reduce the number of condensable species tracked for SOA formation.

The fast-degradation lifetime jumping strategy reduces numerical stiffness by targeting species with very short lifetimes, most commonly arising from rapid unimolecular decomposition (e.g., alkoxy radicals and Criegee biradicals) or fast bimolecular reactions (e.g., reactions with O_2). Because GECKO-A already implicitly bypasses certain short-lived intermediates during mechanism generation, including carbon-centered radicals and excited Criegee intermediates, this strategy primarily targets alkoxy radicals retained in GECKO-A mechanism generation. Species with multiple competing degradation channels that have different temperature dependencies are currently excluded from the strategy, as their product distributions vary significantly with temperature.

3.2.2 Environmentally sensitive strategies

Environmentally sensitive TBR strategies depend on environmental parameters such as T, SZA, RH, pressure, and oxidant or inorganic concentrations. Since these parameters can vary widely across chemical regimes, identical thresholds can result in different reduction outcomes under different conditions. As pointed out in Sect. 3.1.2, GENOA v3 supports the use of multiple scenarios, in which case boundary values of the evaluated strategy parameters are used for TBR threshold comparison. The current environmentally sensitive strategies are summarized below, with their threshold notation:

– Absolute product yield removal: Upper-bound absolute product yields are estimated for species under predefined scenarios, with PVOC yields initialized to one and propagated recursively using relative effective yields (Eq. B1). Species with yields below $Y^{\text{abs}}_{\text{min}}$ are removed. This flux-based, path-integrated metric characterizes relative product distribution rather than species abundance, enabling cross-branch comparison and remaining sensitive to the selected scenarios.

– Effective reaction branching ratio removal: Branching ratios are estimated for reactions from pseudo-unimolecular rates (Sect. 3.1.2), normalized against those of all reactions with the same reactants. Reactions with branching ratios below B_{min} are removed, excluding minor pathways with negligible contributions to product formation.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

– Branch concentration removal: Branch concentrations are estimated as the sum of reference concentrations of all species within the branch. This strategy is applied only to species branches that include downstream species in addition to the root species. Branches with concentrations below $C_{\text{min}}^{\text{ref,brh}}$ are removed.

3.2.3 Application of TBR strategies

TBR strategies are evaluated sequentially in a predefined order, with all reduction candidates that satisfy the specified criteria accepted. The selection of strategies, their application order, and the associated threshold values are user-defined and remain fixed during the reduction. The appropriate choice of these settings may vary across mechanisms to achieve effective size reduction while maintaining accuracy. Some strategies are independent of subsequent changes in the mechanism and therefore require only a single execution (e.g., environmentally insensitive strategies related to P_{sat}), whereas others depend on the evolving mechanism structure and may identify additional reduction candidates after earlier strategies have been applied. As a result, the TBR 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.

As noted above, threshold cutoffs are also applied during the generation of explicit mechanisms, for example in GECKO-A, to limit excessive growth of chemical schemes for large PVOCs (Szopa et al., 2005; Aumont et al., 2005). These are conceptually similar to some TBR strategies, thus when comparable cutoffs have already been applied during mechanism generation, the additional reduction achievable by the corresponding TBR strategies may be limited. The final reduced mechanism is then evaluated using the metrics described in Sect. 3.3.3.

3.3 Simulation-based reduction (SBR)

SBR, referred to as training in GENOA v1 and v2, follows a hierarchical structure in which the concepts of stage, cycle, and step (explained in Sect. B2) support structured configuration and progressive reduction.

Conceptually, an SBR step follows an iterative three-part structure: identification of reduction candidates, evaluation of their performance, and selection of the optimal candidate. In each step, GENOA processes a predefined group of species and associated reactions (Sect. 3.3.1) and generates a set of base and variant candidates (Sect. 3.3.2), where each candidate represents an alternative modification to the same portion of the mechanism. These candidates are evaluated by comparing reduction efficiency and simulation accuracy against a reference mechanism under predefined scenarios, as described in Sect. 3.3.3. The candidate achieving the best balance between reduction efficiency and accuracy 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.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

3.3.1 Species grouping and processing order

GENOA v3 implements a configurable species grouping and ordering scheme during SBR preparation. In contrast to v1 and v2, which process the species list directly, the grouping now serves as the basic unit for identifying reduction candidates, following a defined ranking order (Fig. 2).

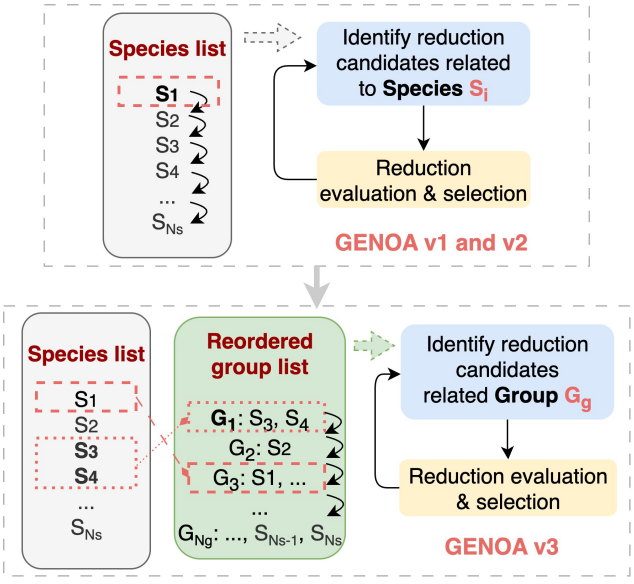

Figure 2. Illustration of the species grouping and ordering scheme in GENOA v3 (lower panel), compared to species-by-species reduction in earlier versions (upper panel). This design supports scalable reduction of large mechanisms through group-based candidate identification and evaluation.

Species are grouped according to shared chemical or physical properties, including molecular formula, volatility, and chemical stability. Each group is then assigned a numerical rank that determines the sequence in which groups are processed during iterative reduction. The ranking prioritizes groups with lower expected impact on reduction accuracy, providing a practical and systematic order for reduction rather than a chemically optimal hierarchy. A detailed mathematical description of the grouping and ranking procedure is provided in Appendix B3, with an illustrative example in Appendix B5.2. Once the order is established, GENOA processes groups sequentially, generating and evaluating reduction candidates on a group-by-group basis.

This grouping-and-ranking supports scalable reduction of very large mechanisms by structuring candidate identification and evaluation at the species group level, substantially improving computational efficiency. The method also supports user-defined grouping criteria and ranking weights. Species-level reduction can still be performed by setting the group size to one, which is generally practical only after partial mechanism reduction. In such cases, GENOA v3 retains the v2 species-ranking method, which assesses species importance using individual removal tests and associated performance metrics (Wang et al., 2023).

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

3.3.2 Identification of base and variant candidates

Based on the reordered group lists, GENOA identifies reduction candidates. For each species group and its associated reactions, GENOA v3 employs reduction strategies that propose modifications to species representation or connectivity without applying fixed thresholds. The SBR strategies are classified into four main types, consistent with previous versions:

- Removal: including the removal of species, reactions, or gas-particle partitioning for condensable species;
- Jumping: skipping intermediate species by connecting precursors directly to products;
- Lumping: merging multiple species into a new surrogate species;
- Replacement: substituting species with chemically similar species already present in the mechanism.

Mathematical formulations and updated implementations of these strategies are provided in Appendix B4, while a step-by-step illustration of base and variant candidate generation is given in Appendix B5.3. Since SBR strategies are applied to the same species group and its associated reactions, the resulting reduction options (termed base candidates) are directly comparable. A base candidate may involve multiple modifications to both species and reactions, and in some cases, partial reductions can yield higher accuracy than base candidates for certain species groups.

To support a more flexible and comprehensive exploration of reduction possibilities, GENOA v3 introduces variant candidates derived from the base candidates. Variants are generated by selecting different combinations of the proposed modifications or by adjusting their parameters. For example, a candidate identified via lumping involving several species may yield variants that lump only a subset of those species, or use different weighting ratios within the same group, resulting in surrogate species with distinct properties. Likewise, a base candidate identified by removal that modifies multiple reactions can generate variants affecting only a subset of those reactions. Together, base candidates and their variants provide a comprehensive set of reduction options. The number of generated variants is controlled to manage computational cost. For large mechanisms, variant generation is typically limited to a small set, while for small mechanisms, a broader range of variant candidates may be explored.

3.3.3 Candidate evaluation criteria and metrics

Base and variant reduction candidates are evaluated and selected based on quantitative metrics for both reduction efficiency and simulation accuracy. Reduction efficiency is assessed using a score reflecting the decrease in mechanism size, specifically in the number of reactions, species, and condensable species. This reduction score (Eq. C3) incorporates user-defined weighting factors to emphasize different aspects of mechanism complexity. Candidates with higher scores indicate greater efficiency.

As the core of SBR, accuracy is evaluated by comparing the temporal concentrations of selected species between the candidate-modified mechanism and the reference mechanism without modifications, using box-model simulations under predefined scenarios. The key configurable elements in the accuracy evaluation are summarized in Table 2 and are listed below:

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

- Evaluation target: Selected species or species groups whose concentrations serve as benchmarks for accuracy. These may include individual compounds, tracers, or aggregated quantities from reducible species, such as tSOA (the default metric for SOA-focused applications).
- Reference mechanism: Baseline mechanism containing evaluation targets of interest and serving as the comparison for evaluating accuracy. This may be either a fixed mechanism defined before reduction or a dynamic reference mechanism updated at each iteration.
- Evaluation scenario: Predefined scenarios used for simulation. These are typically representative of ambient or chamber conditions or tailored to a target environment.
- Error measure: Methods used to quantify deviations between the candidate and reference mechanisms under the evaluation scenarios. The resulting metrics are referred to as reduction errors and are used to guide candidate evaluation.
- Error tolerances: Thresholds applied to reduction errors that determine whether a candidate is acceptable. A candidate is considered valid only if all predefined error tolerances are satisfied.

All listed elements can be specified by the user, with values that can vary across SBR stages, allowing the reduction strategy to evolve as the mechanism is simplified. Mathematical formulations of error measures for individual simulations and across all simulations are provided in Appendices C2 and C3.

Although multiple error measures can be defined, this work adopts the fractional absolute error (FAE) as the primary metric. An FAE-based reduction error (Eqs. C4–C5) is defined as the maximum FAE computed over day 1 and day 2 of the simulation, capturing both rapid initial chemistry and subsequent chemical evolution. For each reduction attempt, reduction errors are computed across the evaluation targets and scenarios, with mean, maximum, and error increase metrics ($\overline{E}$, $E^{\max}$, $\Delta\overline{E}$) compared against corresponding tolerances (ϵ_{avg}, $\epsilon_{\max}$, ϵ_{Δ}) for tolerance control (Appendices C2–C4). This metric is used consistently for SBR candidate evaluation and reduction performance assessment, and is referred to as the reduction error throughout this work.

4 Model evaluation: setup and configuration

4.1 Test mechanism selection

To evaluate the GENOA v3 reduction process, highly detailed VOC mechanisms generated by GECKO-A v1.0 (Aumont et al., 2025b) are used as input mechanisms for reduction. Each mechanism includes up to six generations of oxidation and degradation of a single PVOC from one of three major atmospheric VOC classes (alkanes, alkenes, and monoterpenes). All mechanisms are produced using default GECKO-A parameters at a reference temperature of 298 K, with isomer substitution enabled. Mechanism generation applies cutoff criteria that truncate chemistry for species with P_{sat} below 10^{-13} atm, prune reaction pathways with branching ratios below 5 %, and remove products with effective yields below 0.1 %. Key species

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

properties, including P_{sat} calculated using the method of Nannoolal et al. (2004, 2008), are determined during mechanism generation.

Prior to reduction, input mechanisms undergo a preparation step in which reactions with identical reactants and reaction types are merged when their rate coefficients are proportional. This step removes redundancy while preserving total production and loss rates and is independent of the chemical environment. For large GECKO-A mechanisms, up to approximately 10 % of reactions can be merged during this step. An example illustrating reaction merging is provided in Appendix B5.1. Gas–particle partitioning for highly volatile species ($P_{sat} > 1 \times 10^{-3}$ atm) and for all PVOCs is also removed during preparation, which does not introduce significant differences in the simulation results.

In TBR, both the starting-point and reference mechanisms correspond to the input mechanism after the preparation step. The PVOCs span a broad range of carbon numbers, including n-alkanes (C_nH_{2n+2}, "C5"–"C16"), 1-alkenes (C_nH_{2n}, "U5"–"U16", and monoterpenes ($C_{10}H_{16}$: α-pinene, β-pinene, limonene, camphene, and carene). Here, "C" and "U" are used to distinguish saturated (alkane) and unsaturated (alkene) homologous series, respectively, with the number indicating the PVOC carbon number. When referring to these VOC classes as a whole, they are denoted as "Cs", "Us", and "MTs", respectively.

For SBR, octene ("U8") is selected as the primary example mechanism due to its intermediate size and chemical complexity among all test mechanisms. Its unsaturated chemistry contributes to both gas-phase reactivity and SOA formation, allowing reduction performance to be evaluated against multiple metrics. Reductions of dodecane ("C12") and limonene ("LIM") mechanisms are also included for cross-mechanism comparison. 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 in Sect. 4.4.

4.2 Evaluation scenario specification

A total of 11 scenarios are used to define the environmental parameters, including profiles of oxidants and key inorganic species. The scenarios are designed to span a wide range of chemical regimes relevant to atmospheric mechanism modeling, including idealized box-model environments and observation-based diurnal conditions, and therefore vary in their representativeness of the ambient atmosphere. The scenarios, together with their abbreviated names and key differences, are summarized in Table 3 and are introduced below:

- Five scenarios representing typical atmospheric chemistry in urban, remote, continental, polluted continental, and remote continental environments (urban, remote, continental, poll-cont, and remote-cont), each with fixed O_3 and NO_x concentrations during the simulations. Other oxidants and inorganic species are specified only as initial conditions. These idealized scenarios follow previous studies using the GECKO-A framework (e.g., Camredon et al., 2007), with oxidant and NO_x levels selected to be consistent with configurations from the GECKO-A output library (Lee-Taylor and Drews, 2024).

- One simplified scenario (simplified) in which O_3 and NO_x are specified only as initial concentrations. No additional species are initialized.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

- Four fixed-photolysis scenarios derived from the urban and simplified scenarios, including no-photolysis (urban-nophot, simpl-nophot; SZA at 90 °) and high-photolysis (urban-highphot, simpl-highphot; SZA at 0 °) cases.
- One diurnally varying scenario (diurnal) featuring hourly oxidant, NO_x, and PVOC concentrations to simulate daily variation. This is the only scenario with diurnal PVOC inputs. The diurnal constraints are adopted from example files distributed with the F0AM box model framework (Wolfe et al., 2016, Table S2), based on observations from the 2013 Southern Oxidant and Aerosol Study field campaign (e.g., Kaiser et al., 2016).

Table 3. Summary of the differences among the 11 evaluation scenarios. [a]

Scenario	O_3 (ppb)	NO_x (ppb)	Init+	SZA
urban	fixed (40.6)	fixed (20.3)	Yes	nf
remote	fixed (40.6)	fixed (0.01)	Yes	nf
continental	fixed (40.6)	fixed (0.51)	Yes	nf
poll-cont	fixed (40.6)	fixed (2.03)	Yes	nf
remote-cont	fixed (40.6)	fixed (0.025)	Yes	nf
simplified	init (40.6)	init (20.3)	No	nf
urban-nophot	fixed (40.6)	fixed (20.3)	Yes	90
urban-highphot	fixed (40.6)	fixed (20.3)	Yes	0
simpl-nophot	init (40.6)	init (20.3)	No	90
simpl-highphot	init (40.6)	init (20.3)	No	0
diurnal	diurnal	diurnal	Yes	nf

[a] Columns (left to right): abbreviated scenario names (bold used for SBR candidate evaluation; Sect. 4.4); O_3 and NO_x concentrations (fixed: held constant throughout the simulation; init: only initial values are prescribed); Init+: whether additional species (beyond O_3 and NO_x) are initialized (full list in Tables S1 and S2); SZA: nf denotes diurnally varying; 0 ° and 90 ° denote fixed high- and no-photolysis scenarios.

Non-reducible inorganic reactions are fully retained in all VOC mechanisms, allowing oxidant and hydrocarbon concentrations to evolve according to scenario-specific constraints. Simulations start at midnight (0 h) and run for two full days (48 h), using a uniform time step of 600 s. The GECKO-A box model v1.0 (originally from Aumont et al. (2025a)) is used and modified only at the interface level to enable coupling with GENOA v3, without affecting the core simulation results. T is fixed at 298 K and RH at 70 %. A high initial SOA 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. For all non-diurnal scenarios, the initial PVOC concentration is set to 10 ppb.

The Supplement (Sect. S1) summarizes the detailed scenario inputs (Tables S1–S2; Sect. S1.1) and presents scenario-specific time series of key oxidants and radicals simulated with test mechanisms, organized by VOC class and shown as minimum–maximum ranges and means (Figs. S1–S6; Sect. S1.2), illustrating the diversity of the selected scenarios. While these scenarios span a range of chemical conditions, they are intentionally simplified and not intended to represent any specific or fully realistic

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

atmospheric environment. Instead, they provide a consistent and controlled set of conditions for demonstrating and evaluating the GENOA v3 reduction workflow. Further improvements to scenario design, including multi-precursor mixtures and more realistic oxidant evolution, are important for developing application-ready reduced mechanisms but are beyond the scope of this work and should be explored in future studies.

All scenarios are used for reduction identification in both TBR and SBR through the estimation of reference concentrations and other reduction-related metrics. For SBR candidate evaluation, a representative subset of six scenarios is used to reduce computational cost while maintaining chemical diversity: urban, urban-nophot, urban-highphot, remote, simplified, and diurnal, highlighted in bold in Table 3. Although SBR is guided by these six scenarios, its performance is evaluated using 11 scenarios, with results primarily reported as means unless otherwise specified, consistent with the TBR evaluation in Sect. 5.

4.3 TBR sensitivity analysis and strategy selection

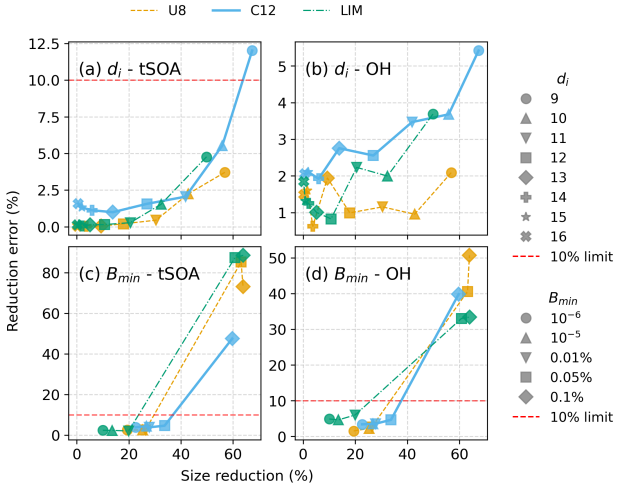

Figure 3. Mechanism size reduction (f_{rdc}; Eq. C2) versus mean reduction error (Eq. C6) in tSOA (left panels) and OH (right panels) for the node depth limit (d_i, top panels) and reaction branching ratio removal (B_{min}, bottom panels) TBR strategies across U8, C12, and LIM mechanisms. Marker shapes denote threshold values, with lines connecting results from the same mechanism. Errors are computed as the mean across all 11 scenarios.

The effectiveness of TBR depends on the choice of reduction strategies and their associated threshold values. To guide threshold selection, sensitivity analyses are performed for individual TBR strategies, examining the trade-off between mechanism size reduction and simulation accuracy as thresholds and evaluation targets are varied. Figure 3 presents the sensitivity of representative TBR strategies, including node-depth limitation and reaction branching-ratio removal, evaluated using tSOA and OH as key SOA-related and gas-phase evaluation targets, respectively. Additional sensitivity results are provided in the Supplement (Sect. S2.1), organized by strategy (Figs. S7–S15), and include tSOA, OH, and additional evaluation targets NO_3 and lOH (OH loss tracer; Sect. 4.4.1).

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

As expected, larger mechanism size reductions are generally achieved at the expense of increased error. TBR strategies differ fundamentally in their sensitivity to threshold values: some allow progressive size reduction with gradual error increases (e.g., Figs. 3(a)–(b), while others lead to sharp error amplification once thresholds exceed narrow ranges (e.g., reaction branching-ratio removal in Figs. 3(c)–(d)). These responses vary across mechanisms and evaluation targets, requiring strategy-specific threshold selection. As noted in Sect. 3.2.3, several TBR strategies overlap conceptually with cutoff criteria already applied during GECKO-A mechanism generation (Sect. 4.1), including those related to NVOCs, low branching-ratio reactions, and pathways with low effective yields. In GECKO-A, such cutoffs are imposed during mechanism construction to limit scheme growth, whereas in GENOA analogous TBR strategies are applied to simplify existing mechanisms. When similar thresholds are used in both steps, the scope for additional TBR is limited. For example, NVOC downstream removal with $P_{\mathrm{sat}}^{\mathrm{nvoc,dwn}} = 1 \times 10^{-13}$ atm produces no reduction because the corresponding pathways are already removed during generation. This effect is considered when selecting the TBR strategy set.

Based on these analyses, suitable threshold values were selected for each strategy that achieve relatively large size reduction while maintaining acceptable accuracy across mechanisms and evaluation targets (errors near 10 %). As effective reduction requires the combined application of multiple strategies, TBR performance is evaluated using two representative strategy sets:

- T6: comprising the six environmentally insensitive strategies;
- T10: comprising all currently implemented TBR strategies, including both environmentally insensitive and environmentally sensitive strategies.

T10 extends T6 to a total of ten strategies, providing greater reduction efficiency at the expense of increased sensitivity to environmental conditions. The selected strategy sets and adopted threshold values are summarized in Table 4. Within each set, strategies are applied sequentially in the order listed in the table, following the procedure described in Sect. 3.2.3.

4.4 SBR configuration

SBR is applied to U8, C12, and LIM mechanisms. Configurations are kept consistent across mechanisms to enable direct comparison, with minor adjustments to account for mechanism-specific sensitivities. The starting-point mechanisms (denoted as StP for SBR) are obtained from the corresponding reference mechanisms (Ref) through TBR with the T10 strategy set. To reduce computational cost during SBR, the original T10 thresholds are slightly loosened to generate StP mechanisms. Specifically, the node depth limit (d_i) is reduced from 12 to 10 for all mechanisms, and for U8 the reaction branching ratio threshold (B_{min}) is increased from 0.01 % to 0.1 %.

4.4.1 Evaluation targets and tracers

Two evaluation configurations are applied: an SOA-focused reduction evaluates performance using tSOA only to prioritize SOA formation, while a multi-target reduction extends the evaluation to include oxidants and inorganic tracers to capture broader impacts on atmospheric chemistry. For multi-target configurations, tracer species are used to quantify the production and loss of key oxidant and inorganic species, including OH, NO_3, O_3, NO_2, and HO_2, associated with reactions in which

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

Table 4. TBR strategies and threshold values used in T6 and T10 strategy sets.

Reduction strategy	Threshold	Value
Reduction strategies for T6		
Fast-degrading species jumping	τ^{fast}	1 s
Node depth limit	d_i	12
NVOC reclassification	$P_{\text{sat}}^{\text{nvoc}}$	1×10^{-10} atm
NVOC downstream removal	$P_{\text{sat}}^{\text{nvoc,dwn}}$	1×10^{-10} atm
VOC reclassification	$P_{\text{sat}}^{\text{svoc}}$	1×10^{-5} atm
VOC branch removal	$P_{\text{sat}}^{\text{svoc,brh}}$	1×10^{-5} atm
Additional strategies for T10 [a]		
Yield-to-depth ratio removal	$\left(\frac{Y^{\text{rel}}}{d}\right)_{\text{min}}$	0.01
Branch concentration removal	$C_{\text{min}}^{\text{ref,brh}}$	0.1 ppt
Absolute product yield removal	$Y_{\text{min}}^{\text{abs}}$	1×10^{-8}
Reaction branching ratio removal	B_{min}	0.01 %

[a] T10 includes all strategies in T6 plus four additional strategies. For SBR starting-point mechanism preparation (Sect. 4.4), a TBR step using T10 is applied, with d_i and B_{min} adjusted for more effective reduction.

reducible species are formed or consumed. These tracers are diagnostic quantities that do not participate in the chemistry but integrate reaction fluxes from reducible pathways. Table B5 provides an example of tracer implementation in a simplified mechanism.

In practice, tracers enable quantitative evaluation in scenarios where evaluation targets follow fixed profiles or where target concentrations, such as radicals, are too low to serve as reliable error metrics. By accumulating over time, tracer concentrations provide a practical measure for assessing whether essential chemical processes are preserved after reduction. Tracers are generally correlated with their corresponding species, but the strength of this relationship depends on the underlying chemistry and on whether changes in species concentration are mainly driven by production or loss pathways represented by the tracer. Consequently, tracer–species correlations vary across species, mechanisms, and scenarios. All production and loss tracers with non-negligible concentrations simulated using the reference mechanism under the evaluation scenarios are selected as evaluation targets. The resulting set includes production and loss tracers for NO_2 and HO_2 and loss tracers for OH, NO_3, and

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

O_3. No O_3 tracer is included for C12. For each mechanism, the evaluation target sets (S^{eval}) are defined as:

$$S^{\text{eval}}_{\text{U8A}} = \{\text{tSOA}\},$$
$$S^{\text{eval}}_{\text{U8B}} = \{\text{tSOA}, \text{lO}_3, \text{lNO}_3, \text{lOH}, \text{pNO}_2, \text{lNO}_2, \text{pHO}_2, \text{lHO}_2\},$$

$$S^{\text{eval}}_{\text{C12A}} = \{\text{tSOA}\},$$
$$S^{\text{eval}}_{\text{C12B}} = \{\text{tSOA}, \text{lNO}_3, \text{lOH}, \text{pNO}_2, \text{lNO}_2, \text{pHO}_2, \text{lHO}_2\},$$

$$S^{\text{eval}}_{\text{LIMA}} = \{\text{tSOA}\},$$
$$S^{\text{eval}}_{\text{LIMB}} = \{\text{tSOA}, \text{lO}_3, \text{lNO}_3, \text{lOH}, \text{pNO}_2, \text{lNO}_2, \text{pHO}_2, \text{lHO}_2\}.$$

Each evaluation target set is indexed by mechanism and configuration identifier ("A" for SOA-focused, "B" for multi-target). Tracer names are constructed from the corresponding species, prefixed with "p" or "l" to denote production or loss within the mechanism.

Currently, tracer selection is based on their concentrations. Tracer–species response consistency is assessed preliminarily using TBR sensitivity tests by comparing reduction error correlations across reduced mechanisms. Figure S16 in the Supplement shows strong correlations (> 0.7) for OH and HO_2 tracers, whereas correlations for NO_x, NO_3, and O_3 are weaker and vary across VOC classes. These tracer–species relationships are further discussed in the evaluation of SBR reduction performance (Sect. 6).

4.4.2 Stage configuration and error tolerances

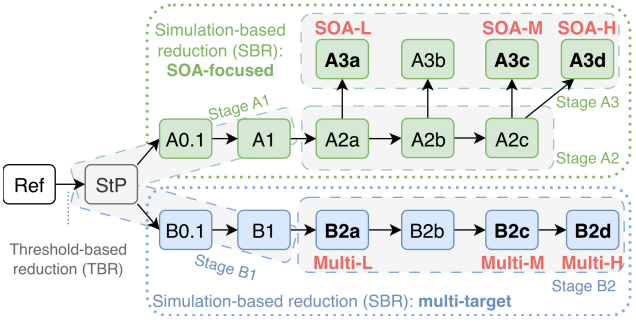

Figure 4. Illustration of the SBR workflow. The SOA-focused reduction (upper, green) proceeds through three stages, while the multi-target reduction (lower, blue) proceeds through two stages. Mechanisms are named according to Table 5.

SBR is performed in sequential stages (Table 5; Fig. 4). SOA-focused reduction uses three stages (A1–A3), while multi-target reduction uses two (B1 and B2). Both reductions start from the same starting-point mechanism (StP), with errors evaluated relative to the reference mechanism (Ref) using the metrics described in Appendices C2–C5. Each stage includes multiple re-

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

Table 5. Stage configurations and mechanisms for SBR evaluation.

Stage	Mechanism [a]		Error tolerance (%) [b]		
	Start	End	ϵ_{max}	ϵ_{avg}	ϵ_{Δ}
SOA-focused reduction (A)					
A1	StP [c]	A0.1, A1 [d]	15	5	0.5
A2	A1	–	15	5	0.5, 5 [e]
	–	A2a	15	5	5
	–	A2b	20,30	8,10	8,10 [f]
	–	A2c	50	15	15
A3-1	A2a	**A3a (SOA-L)** [g]	15	5	5
A3-2	A2b	A3b	30	10	10
A3-3	A2c	**A3c (SOA-M)**	30	10	10
A3-4	A2c	**A3d (SOA-H)**	50	15	15
Multi-target reduction (B)					
B1	StP	B0.1, B1 [d]	35	5	0.5
B2	B1	–	15	5	0.5
	–	**B2a (multi-L)**	15	5	5
	–	B2b	20	8	8
	–	**B2c (multi-M)**	30	10	10
	–	**B2d (multi-H)**	50	15	15

[a] Start and End denote the input and output mechanisms of each stage; "–" indicates continuation within a stage without a new mechanism identifier.
[b] Error calculation and tolerance definitions in Appendices C2–C4. Values shown are primarily for U8; for C12 and LIM, Stages A1 and B1 use 35 / 10 / 1, and the first ϵ_{Δ} in Stages A2 and B2 is 1. [c] StP: starting point mechanism for SBR. [d] A0.1 and B0.1 denote mechanisms after the first reduction cycle in Stages A1 and B1. [e] First set to 0.5 %, then 5 %, while other tolerances remain unchanged.
[f] Two tolerance sets are used: (20, 8, 8) followed by (30, 10, 10).
[g] Final reduced mechanisms with low (-L), medium (-M), and high (-H) error tolerances are highlighted in bold and assigned secondary names.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

duction cycles (Appendix B2), with stage-specific settings defined as lists and applied sequentially across cycles (e.g., different error tolerances). Stages A1 and B1 perform the initial reduction from StP and are configured for efficient size reduction of large mechanisms. Mechanisms A0.1 and B0.1 denote the intermediate mechanisms obtained after the first reduction cycle in Stages A1 and B1, respectively, and are recorded to illustrate early-stage reduction behavior. In later stages (A2, A3, and B2), variant candidates are enabled to explore additional reduction possibilities as the mechanism size decreases. Stage transitions follow the mechanism outputs defined in Table 5, which correspond to reduced mechanisms from the previous stage after completing reduction cycles under specified error tolerances. These mechanisms are generally reorganized for the next stage, with reactions merged by reactants and kinetic types. An exception is the early-stage C12 reduction (A1, A2, and B1), where merging is not applied due to the complexity of the resulting reactions.

For SOA-focused reduction, Stage A2 introduces an aerosol-oriented constraint to prioritize removal of condensable species. Only candidates that reduce the reduction errors or decrease condensable species counts are accepted. Stage A3 removes this constraint and applies error-based selection only, using tolerances comparable to or slightly relaxed from those in Stage A2. As a result, Stage A3 produces four branches (A3-1–A3-4) originating from different A2 mechanisms, rather than following a single sequence as in Stages A2 and B2. For multi-target reduction, Stage B2 applies error tolerance settings similar to those used in Stages A2 and A3, without the aerosol-specific constraint, resulting in reduced mechanisms (B2a–B2d) analogous in accuracy to A3 mechanisms. Additional details are provided in the Supplement (Sect. S3.1), including size and accuracy evolution from the StP to A3d mechanisms during U8 reduction (Fig. S20).

For each reduction, three representative final reduced mechanisms are highlighted and assigned a secondary naming scheme, denoted by the suffixes -L, -M, and -H, corresponding to low, medium, and high error tolerance levels. Each suffix is combined with the reduction configuration, using “SOA-” for SOA-focused reductions and “multi-” for multi-target reductions (Table 5). These tolerances target mean reduction errors of approximately 5, 10, and 15 % (with maximum errors of 15, 30, and 50 %), although the actual achieved errors may differ slightly. These labeled mechanisms correspond to A3a / A3c / A3d and B2a / B2c / B2d, respectively, as highlighted in Table 5 and Fig. 4. This secondary naming provides a clear and consistent way to compare mechanisms with similar accuracy across reductions and configurations. When original SBR identifiers are presented, these mechanisms are highlighted in bold hereafter in tables and figures.

4.5 Evaluation metrics and analysis conventions

TBR and SBR performance is presented and discussed in Sects. 5 and 6, respectively. The evaluation focuses on characterizing reduction behavior and its relation to chemically interpretable differences among mechanisms. Application-ready mechanism development will be the focus of future work. Reduction performance is quantified primarily using the reduction error (Eq. C6). The normalized mean bias (NMB), reported with the mean values, serves as a complementary diagnostic of systematic deviations. By default, results are evaluated across the full set of evaluation scenarios described in Sect. 4.2 (11 scenarios for general evaluation and six scenarios for SBR candidate evaluation). Scenario-specific results are discussed selectively when they reveal systematic or chemically interpretable behavior. A limited set of representative metrics is emphasized in the main paper, while additional supporting metrics are provided in the Supplement (Sects. S2–S3).

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

5 Evaluation of TBR performances

TBR performance is evaluated using the size–accuracy trade-off of reduced mechanisms generated with the T6 and T10 strategy sets, with size reduction quantified by f_{rdc} (Sect. 5.1) and accuracy assessed using SOA-related (Sect. 5.2) and gas-phase (Sect. 5.3) metrics. Detailed statistics for size-related metrics are provided in Table S3, while accuracy-assessment metrics across PVOCs and scenarios are provided in Tables S4–S5.

5.1 Size reduction

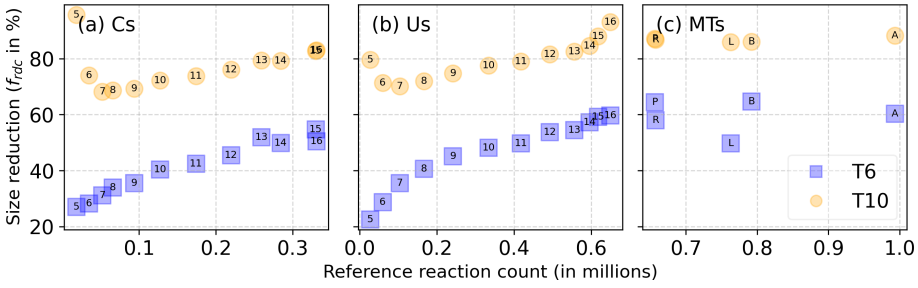

Figure 5. Size reduction (f_{rdc}; Eq. C2) of reduced TBR mechanisms as a function of the number of reactions in the reference mechanism. Each symbol denotes a reduced mechanism derived from a reference PVOC degradation mechanism. Marker labels indicate PVOC identity by carbon number for Cs and Us and by compound name for MTs. The applied TBR strategy set is indicated by the symbol shape, with squares denoting T6 and circles denoting T10. For MTs, marker labels denote individual compounds: A = API (α-pinene), B = BPI (β-pinene), L = LIM (limonene), R = CRN (carene), and P = CMP (camphene).

Figure 5 summarizes TBR size reductions across VOC classes, where larger values correspond to greater size reduction. Reference mechanism size varies strongly with PVOC size and structure. For Cs and Us, reaction counts span approximately three orders of magnitude (10^3–10^6), with Cs consistently smaller than Us at the same carbon number. MT mechanisms are larger than the other ten-carbon mechanisms (C10 and U10), reflecting higher structural complexity and oxidation complexity.

Across all VOC classes, T6 produces size reductions of 23–65 % (46 % mean), while the more aggressive T10 increases reductions to 68–96 % (80 % mean), nearly doubling the reduction through the inclusion of four environmentally sensitive TBR strategies. For Cs and Us, size reduction increases with reference mechanism size and PVOC carbon number, consistent with larger mechanisms containing more low-yield or redundant pathways that are amenable to reduction. Deviations from this trend occur for the smallest mechanisms (C5–C6 and U5–U6) under T10, suggesting potential over-reduction of smaller reference mechanisms with environmentally sensitive strategies. Among the ten-carbon systems, MTs exhibit the largest size reductions, followed by U10 and C10. Differences among MTs are small (50–65 % for T6 and 86–88 % for T10) relative to variations across Cs and Us with carbon number, consistent with a secondary influence of isomer structure.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

5.2 SOA formation and composition

SOA formation is evaluated using tSOA as the primary metric, while SOA composition is evaluated using bulk elemental ratios and particle-phase functional group distributions. Evaluation of SOA-related metrics excludes C5–C6 from the main discussion due to negligible SOA formation, with mean tSOA of 0.02 and 0.12 μg m^{-3} (SOA yields of 0.2 % and 0.6 %). These mechanisms are retained for gas-phase metric evaluation.

5.2.1 Total SOA (tSOA)

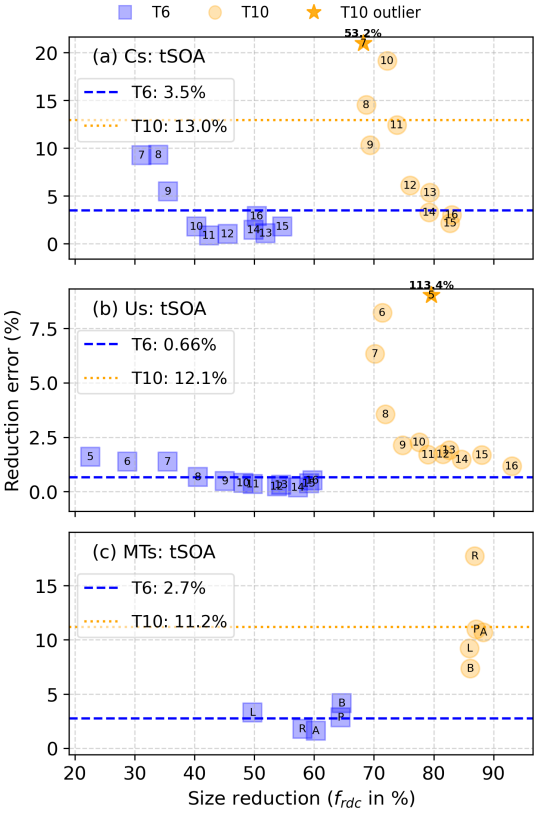

Figure 6. Size reduction versus mean tSOA reduction error across 11 scenarios for TBR mechanisms. Symbols follow Fig. 5. Horizontal lines show overall mean errors for T6 (blue, dashed) and T10 (orange, dotted). C5–C6 are excluded due to negligible SOA production. T10 outliers (C7 and U5) are marked by stars and annotated with their actual errors.

Figure 6 illustrates the size–accuracy trade-offs between mechanism reduction and tSOA preservation across VOC classes. Across all mechanisms, tSOA is generally well preserved. T6 yields low mean and median tSOA errors (2.1 % and 0.73 %), while T10 increases these values to 12.3 % and 4.2 %, remaining acceptable given the corresponding size reduction fractions. The mean exceeds the median, particularly for T10, reflecting a small number of large-error reductions. tSOA errors vary

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

with VOC class and mechanism size. Us show the smallest errors (median 0.28 % and 1.2 % for T6 and T10), followed by Cs (1.5 % and 5.5 %) and MTs (2.3 % and 9.5 %). For Cs and Us, errors generally decrease with increasing PVOC carbon number, particularly under T10. The largest T10 errors occur for the smallest mechanisms, where C7 and U5 are identified as extreme outliers. This error–size relationship is not observed for MTs, for which size reduction and tSOA error show no consistent correlation.

When examined alongside the mean tSOA distribution (Fig. 7(a)), these results suggest a systematic dependence of reduction performance on mechanism structure linked to PVOC carbon number, with improved preservation of tSOA for higher carbon-number PVOCs within the same functional class. As the carbon number increases, early-generation oxidation products form lower-volatility species with larger SOA contributions. Consequently, SOA formation becomes increasingly dominated by early-generation products and a smaller number of high-contribution species, which are more readily retained during reduction for Cs and Us. In contrast, MTs show weaker and less systematic scaling with mechanism size, consistent with structural isomerism and more diverse oxidation pathways.

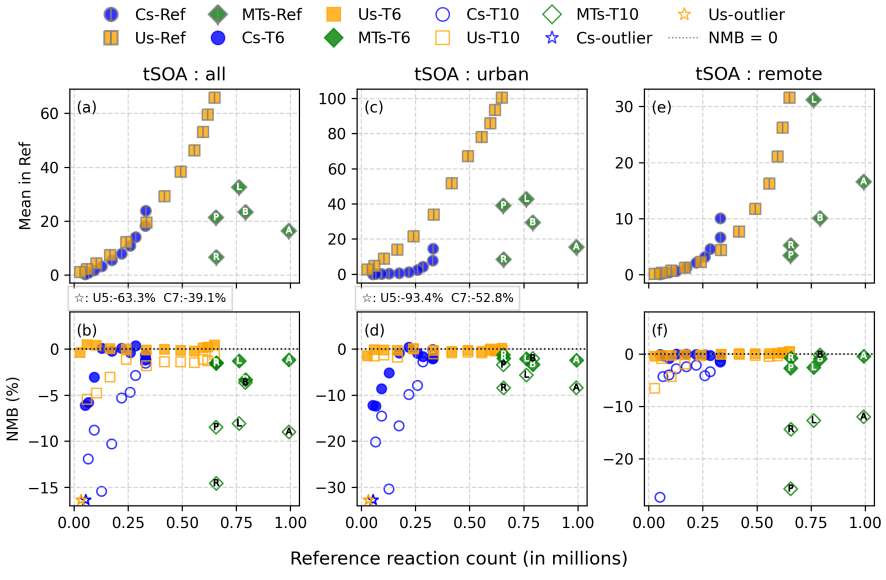

Figure 7. Mean tSOA and NMB of TBR mechanisms relative to the reference. Top panels show mean tSOA (μg m^{-3}) simulated by the reference mechanisms (Ref), and bottom panels show NMB of the TBR mechanisms relative to Ref. Results are shown as means across all scenarios (left), urban (middle), and remote (right) scenario only. C5–C6 data points are excluded. T10 outliers (if any) are marked by empty stars, with PVOC identifiers and bias values listed above the bias panels in the same left-to-right order as the stars along the x-axis. Marker labels are shown for MTs only; for Cs and Us they are omitted to avoid crowding and follow the near-linear relation with the x-axis (e.g., Fig. 5).

Bias diagnostics further help interpret these results, showing predominantly negative NMB (Fig. 7(b)), with larger absolute biases under T10 than T6 (median -3.4 % compared to -0.27 %), indicating systematic tSOA underestimation due to TBR. This

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

is expected, as most TBR strategies act through species or pathway removal, with the exception of lifetime jumping strategy. Under T6, biases remain close to zero for Us (NMB range: -0.40–0.44 %) and are small for MTs (-3.3 – -1.2 %) and Cs (-6.1 – -0.4 %), indicating that the dominant SOA-forming pathways across VOC classes are well represented under environmentally insensitive strategies. Under T10, extreme outliers and systematic underestimation emerge across all VOC classes, reflecting over-reduction caused by cumulative removal of SOA-forming species and pathways under the environmentally sensitive thresholds introduced in T10.

T10 underestimation varies systematically across scenarios with distinct responses under urban and remote NO_x regimes (Figs. 7(b)–(f); additional scenarios in Fig. S17). Under the urban scenario, T10 bias is dominated by Cs (median -12.2 %) and U5 (-93.4 %), whereas under the remote scenario, the largest biases shift to MTs, with CMP reaching -25.8 % NMB (compared to -3.5 % under urban). These differences are consistent with NO_x sensitivities across VOC classes. Alkane SOA formation is strongly NO_x dependent, and under low-NO_x regimes (e.g., remote) dominant SOA-forming pathways are suppressed, such that their reduction under T10 leads to smaller biases for Cs. In contrast, monoterpene SOA formation under low-NO_x regimes is dominated by multifunctional oxidation pathways, resulting in larger T10 bias for MTs under the remote scenario. Persistent T10 outliers for C7 (-52.8 % and -27.4 % under urban and remote) and U5 (-93.4 % and -6.6 %) are also among the largest biases observed under T6, although with substantially smaller magnitudes. These outliers are consistent with short-chain chemistry favoring fragmentation, particularly for U5 due to its terminal double bond, and with stronger reliance on multi-generation products that are partially removed under T6 and more extensively removed under T10. Together, these results show that pathways classified as reducible using scenario-mean metrics can become important contributors in individual scenarios, leading to scenario-dependent TBR performance.

5.2.2 Elemental SOA ratios

Bulk SOA composition is evaluated using elemental ratios, including OM:OC (organic mass to organic carbon), O:C (oxygen to carbon), H:C (hydrogen to carbon), and N:C (nitrogen to carbon). Overall, T6 yields small metric-level mean errors of 0.3–3.4 %, while T10 increases these errors to 0.7–6.6 %. When evaluated across mechanism- and scenario-level means, 95 % of the T6 simulations are with mean errors below 6.8 %, compared to 10.2 % for T10. O:C and N:C, which directly reflect bulk oxidation state and organic nitrate content, show the strongest sensitivity to reduction and are therefore used as diagnostic metrics (Fig. 8(a)–(d); OM:OC and H:C distributions are shown in Fig. S18). In the reference simulations, O:C and N:C decrease with increasing PVOC size for Cs and Us, reflecting decreasing importance of functionalization with increasing carbon number. Among ten-carbon mechanisms, the mean O:C and N:C ratios decrease from MTs (0.82 and 0.12) to U10 (0.55 and 0.09) and C10 (0.45 and 0.05), consistent with greater structural diversity and enhanced multifunctional oxidation in MTs, and a smaller enhancement in U10 relative to C10.

Under T6, Cs show a small positive mean O:C bias (2.5 %), consistent with preferential reduction of numerous low-yield SOA pathways that contribute relatively less oxygen per unit carbon. Under T10, this bias decreases to 1.4 %, indicating partial compensation as additional strategies remove a broader set of SOA-forming pathways, including those with more highly oxygenated SVOCs. This downward shift suggests interactions among multiple strategies when applied together rather than

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

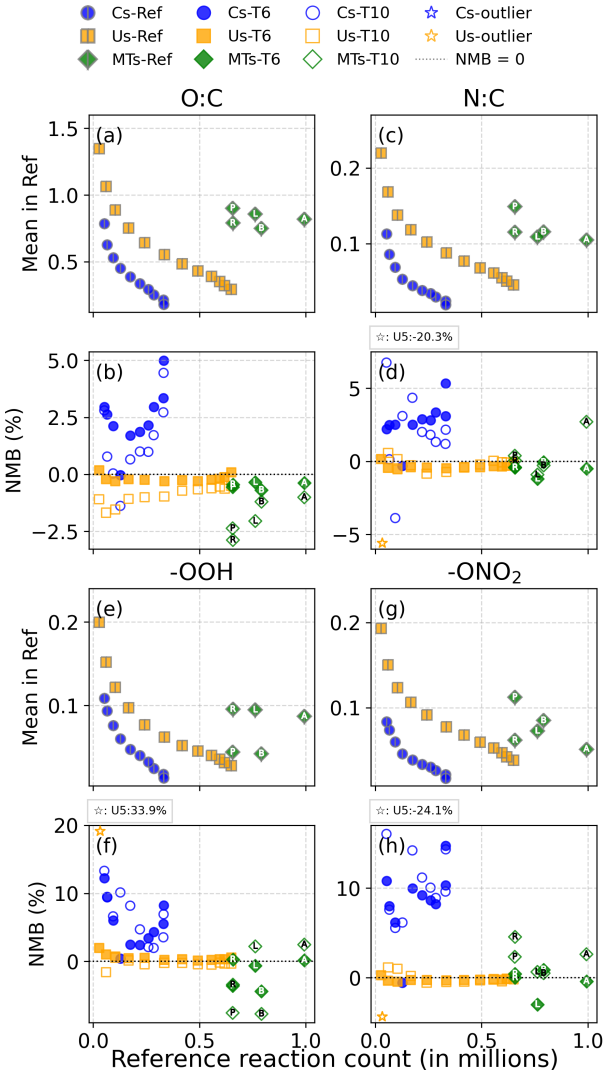

Figure 8. Mean values and NMB of selected elemental SOA ratios and carbon-normalized functional group fractions, formatted as in Fig. 7. Shown metrics include O:C (a–b), N:C (c–d), -ONO_2 (e–f), and -OOH (g–h), with values shown as means across all scenarios.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

a systematic change in oxidation chemistry. Within the Cs series, the non-monotonic O:C NMB pattern, decreasing from C7 (3.0 %) to a near-zero bias at C10 (-0.03 %) and increasing again toward C16 (5.0 %), suggests the interplay of multiple effects. Smaller mechanisms are more sensitive to pathway removal, while intermediate cases show partial compensation between retained and removed pathways. For larger Cs, increased pathway diversity may shift this balance again under fixed TBR thresholds, leading to increasing positive bias. For Us and MTs, mean O:C NMB is slightly negative under T6 (-0.18 % for Us and -0.48 % for MTs), consistent with SOA formation being weighted toward a smaller number of dominant products with higher functionalization compared to Cs. Under T10, the bias becomes more negative (-0.92 % and -1.9 %), consistent with the compensating effect observed for Cs. N:C biases remain small for most mechanisms (median NMB 0.06 % in T6 and 0.29 % in T10), with similar within-PVOC-group distributions to those of O:C, indicating that nitrate-containing SOA is largely preserved. Notable outliers include negative bias for U5 (-20.3 %), consistent with the over-reduction discussed for tSOA.

5.2.3 Functional group composition

SOA composition is also evaluated using carbon-normalized functional group fractions. Two abundant and chemically diagnostic functional groups, organic nitrate (-ONO_2) and hydroperoxide (-OOH), associated with NO_x-dependent and multifunctional oxidation termination, are highlighted (Figs. 8(e)–(h); ketone (>CO) and alcohol (-OH) functional groups are shown in Fig. S18 for reference). Under T6, mean errors are 4.4 % (median 1.4 %) for -OOH and 4.8 % (0.78 %) for -ONO_2, increasing to 8.0 % for both under T10 (medians 3.6 % and 2.7 %, respectively). Bias distributions align with those observed for O:C and N:C, reflecting the direct elemental contributions of -OOH and -ONO_2 to particle-phase SOA composition. Cs show systematic positive bias (median NMB of 2.8 % for -OOH and 2.5 % for -ONO_2), while Us exhibit the smallest bias (0.07 % and 0.01 %), except for U5. MTs display wider case-to-case variability (-0.66 % and 0.45 %), associated with their more complex RO_2 chemistry. These patterns follow the class-dependent trends discussed for tSOA and elemental ratios. Together, elemental ratio and functional-group evaluations show that bulk SOA composition is largely preserved under TBR.

5.3 Gas-phase chemistry

Gas-phase chemistry is evaluated using diverse complementary metrics. Bulk reactivities with OH, O_3, and NO_3 (OHR, O_3R, and NO_3R) are emphasized as key diagnostics, providing integrated measures of gas-phase chemical activity from precursors and oxidation products. Complementary metrics include bulk reactivities of reducible species only (OHR-r, O_3R-r, and NO_3R-r) and concentrations of key oxidants and radicals (NO_x, O_3, NO_3, HO_2, OH, RO_2), as provided in Tables S4–S5 and Fig. S19 to support the interpretation of gas-phase responses. Aggregated across scenarios, T6 largely preserves gas-phase chemistry, with metric-level mean errors of 2.6–7.5 % (mean 4.2 %, median 1.1 %) across all 12 metrics, and 95 % of mean errors evaluated at the individual mechanism–scenario level are below 18.7 %. T10 shows substantially larger deviations, with metric-level mean errors of 7.2–27.3 %, (mean 15.2 % and median 4.2 %), and 95 % of simulations yield mean errors up to 78.6 %, reflecting reduced preservation of gas-phase chemistry. Across VOC classes, MTs show larger errors (median 8.1 %), followed by Us and Cs, which show similar error levels (1.8 % and 1.5 %).

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

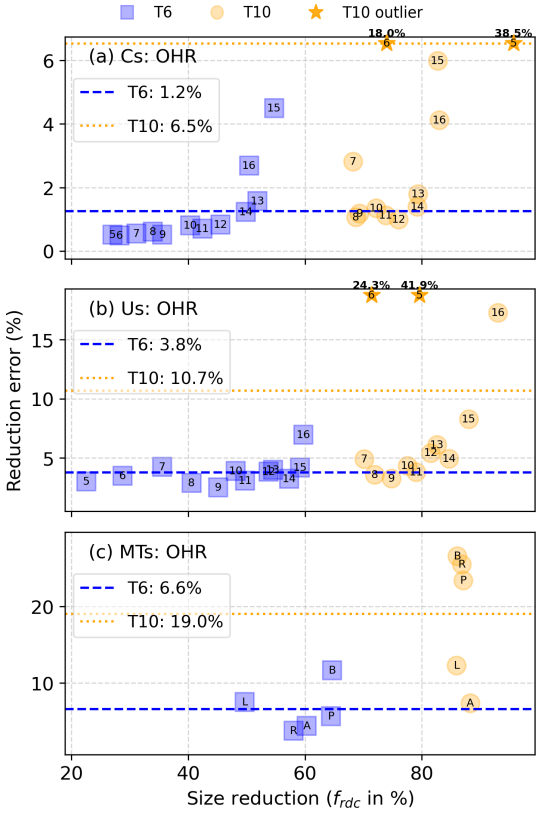

Figure 9. Size reduction versus mean OHR reduction error across 11 scenarios for TBR mechanisms, formatted as in Fig. 6. T10 outliers (C5–C6 and U5–U6) are marked by stars and annotated with their actual errors.

5.3.1 OH reactivity (OHR)

OHR is selected as the primary gas-phase metric, as it reflects the integrated OH sink and overall oxidation capacity from the mechanism. The error distribution follows the overall gas-phase statistics, with smaller errors under T6 (mean 3.2 %) compared to T10 (10.4 %), and class-level errors increasing from Cs (3.9 %) to Us (7.2 %) and MTs (12.8 %), consistent with increasing chemical complexity and a broader distribution of OH-reactive oxidation products. For Cs and Us, the smallest mechanisms (C5–C6 and U5–U6) dominate the upper tail of the T10 error distribution (Figs. 9(a)–(b)), consistent with patterns observed in other gas-phase and SOA-related metrics (e.g., tSOA in Figs. 6(a)–(b)). This result suggests systematic over-reduction under the selected environmentally sensitive TBR strategies and threshold settings in T10, reflected in large size reduction and reduced accuracy for gas-phase and SOA-related metrics. In particular, the current T10 configuration may exclude short-chain oxidation pathways, including C2 and C3 chemistry, that are weakly relevant for SOA formation but substantially contribute to VOC–NO_x–HO_x gas-phase chemistry. The exclusion of these pathways reduces chemical representation in mechanisms

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

dominated by C5–C6 precursors and is consistent with the larger errors observed, suggesting that further refinement of the T10 strategy is needed for these mechanisms when targeted.

Excluding these short-chain outliers (PVOCs with C < 7), OHR errors for Cs and Us remain tightly clustered and largely independent of size reduction, with error showing only a small increase as size reduction increases, distinct from the stronger size dependence observed for tSOA. Relatively larger errors also become more evident for the largest mechanisms (PVOCs with C > 14), including C15–C16 in Fig. 9(a) and U15–U16 in Fig. 9(b). For MTs, errors are similar across isomers under T6 and show increased variability under T10, with BPI exhibiting the largest errors (11.7 % for T6 and 26.6 % for T10; Fig. 9(c)).

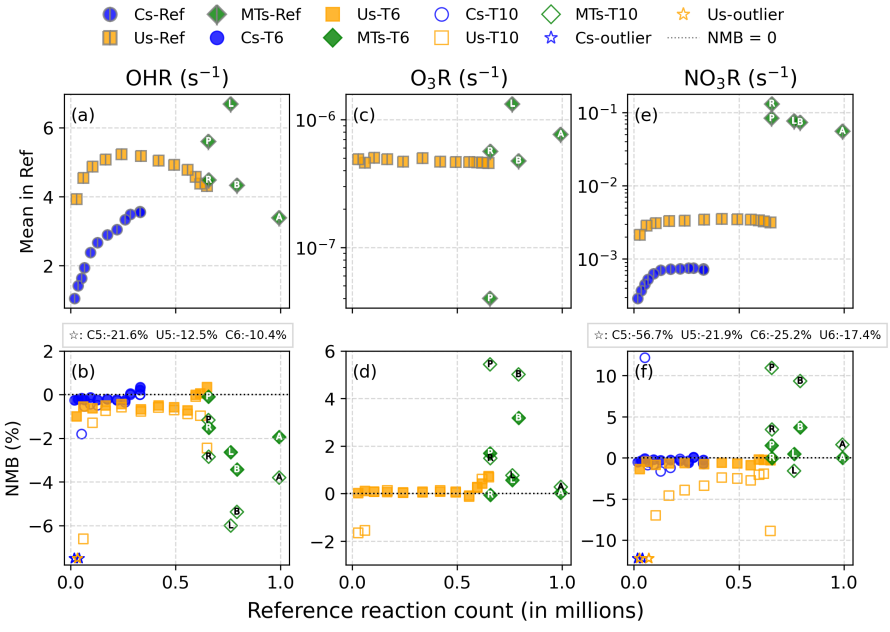

Figure 10. Mean values and NMB of gas-phase bulk reactivities, formatted as in Fig. 7. Panels from left to right correspond to OHR, O_3R, and NO_3R.

Mean OHR in the reference mechanisms shows clear class-dependent patterns (Fig. 10(a)). At comparable carbon numbers, Cs exhibit the lowest OHR (mean 2.6 s^{-1}), followed by Us and MTs (4.7 and 4.9 s^{-1}), reflecting increasing chemical complexity from alkanes to 1-alkenes and monoterpenes. For Cs, OHR increases monotonically with mechanism size, together with increasing reducible-only OHR (OHR-r; Fig. S19 (a)–(b)), consistent with growing contributions from oxidation products in addition to PVOC reactivity as carbon number increases. Us show a non-monotonic dependence of OHR on carbon number, reflecting a balance between competing production and loss processes. OHR increases from U5 to U9, similar to Cs, and decreases for larger Us as gas-to-particle partitioning of lower-volatility products increasingly reduces gas-phase contributions. MTs show larger scatter without systematic size dependence, reflecting their higher intrinsic reactivity and structural diversity. Reflecting the reference OHR patterns (Fig. 10(a)), OHR is generally underestimated under both T6 and T10, with small and tightly clustered biases for Cs and Us (median NMBs of -0.23 % and -0.65 %). Under T10, negative bias becomes more pro-

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

nounced for larger PVOCs (e.g., U15–U16), while the largest deviations (NMB > 10 %) occur for the smallest mechanisms (C5–C6 and U5). MTs show larger biases (median NMB of -2.74 %) than Cs and Us, consistent with the greater scatter in the reference OHR.

5.3.2 O_3R and NO_3R

In addition to OHR, O_3R and NO_3R are evaluated to assess the preservation of unsaturated sinks and nighttime oxidation reactivity (Figs. 10(c)–(f)). O_3R for Cs is excluded from discussion because alkane oxidation in GECKO-A produces negligible unsaturated species, resulting in very small O_3 reactivity (2×10^{-14}–8.3×10^{-9} s^{-1}). For Us, O_3R remains relatively uniform across carbon number, consistent with dominance by the parent alkene double bond and early-generation unsaturated products. NO_3R is consistently high across VOC classes due to the simulations being started with high PVOC mixing ratios (10 ppb) at nighttime conditions. Although the resulting NO_3R levels exceed typical atmospheric values, they are internally consistent and therefore provide a consistent basis for comparing reduction performance. NO_3R increases from Cs to Us and MTs, while remaining tightly clustered across mechanism size within each class.

Similar to OHR, both O_3R and NO_3R are relatively well preserved under T6, with greater scatter and a small number of outliers, consistent with the patterns observed for OHR. For O_3R, Us remain close to zero across carbon number, with slightly positive bias for larger mechanisms (U13–U16) within 0.7 %. Under T10, U5 and U6 show slightly more negative bias (-1.6 % and -1.5 %) compared to other Us. For NO_3R, biases for Cs and Us are predominantly negative, with pronounced outliers for C5–C6 and U5–U6, where T10 biases exceed -15 %. For MTs, by comparison, O_3R and NO_3R exhibit scattered positive biases (1.1 % under T6), which increase under T10 to 2.6 % and 4.7 %, respectively.

Overall, gas-phase reactivity metrics (Figs. 10 and S19) 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 T6, deviations in both gas-phase and SOA-related metrics remain small, whereas under T10 larger deviations emerge, reflecting increased sensitivity to over-reduction and to the cumulative removal of distributed pathways. The metric evaluation suggests that the current uniform, environmentally sensitive strategy 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.

6 Evaluation of SBR performance

SBR performance for U8, C12, and LIM reductions is first evaluated across stages, tracking changes in mechanism size, accuracy, and runtime (Sect. 6.1). Selected SBR mechanisms spanning different accuracy levels and evaluation targets are then assessed based on preservation of key metrics (Sect. 6.2), temporal variability (Sect. 6.3), and dominant reaction pathways (Sect. 6.4).

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

Table 6. Size reduction and accuracy of SBR for U8, C12, and LIM mechanisms. [a]

Mech. [b]	U8 mechanisms		C12 mechanisms		LIM mechanisms	
	N^{rn} / N^{sp} / N^{cd}	$\overline{E}$ / E^{max}	N^{rn} / N^{sp} / N^{cd}	$\overline{E}$ / E^{max}	N^{rn} / N^{sp} / N^{cd}	$\overline{E}$ / E^{max}
Ref	166 821 / 27 221 / 6 836	– / –	219 743 / 34 865 / 7 775	– / –	762 747 / 125 314 / 29 071	– / –
StP [c]	32 662 / 6 770 / 2 215	4.7 / 15.1	40 104 / 7 181 / 2 236	9.0 / 16.4	98 794 / 14 417 / 5 021	9.4 / 15.3
SOA-focused reductions (A series)						
A0.1	616 / 176 / 50	3.0 / 5.7	6 022 (4 857) / 1 203 / 281	6.8 / 8.2	338 / 77 / 22	7.7 / 12.0
A1	274 / 69 / 20	2.6 / 3.9	1 156 (405) / 95 / 65	7.6 / 9.9	294 / 56 / 13	7.0 / 11.0
A2a	207 / 47 / 11	2.1 / 3.3	925 (298) / 89 / 58	4.4 / 5.7	238 / 53 / 12	5.4 / 7.8
A2b	117 / 38 / 6	6.5 / 12.5	637 (142) / 33 / 6	5.3 / 7.3	196 / 47 / 8	5.7 / 8.3
A2c	112 / 37 / 5	10.9 / 16.1	633 (142) / 33 / 5	7.2 / 10.1	97 / 32 / 5	9.2 / 16.4
A3a	52 / 29 / 11	4.9 / 9.1	42 / 20 / 7	5.9 / 12.5	61 / 37 / 9	5.0 / 9.4
A3b	30 / 20 / 6	9.1 / 14.8	31 / 17 / 5	10.9 / 16.7	33 / 23 / 8	9.8 / 17.3
A3c	31 / 19 /5	9.6 / 18.8	37 / 18 / 5	10.2 / 19.0	29 / 19 / 5	9.3 / 19.4
A3d	27 / 15 / 3	14.6 / 28.7	27 / 17 / 5	14.8 / 31.9	27 / 20 / 4	14.5 / 31.3
Multi-target reductions (B series)						
B0.1	1 093 / 268 / 33	5.0 / 12.4	5 696 (4 020) / 846 / 145	8.7 / 19.8	2 160 / 409 / 80	10.0 / 26.7
B1	628 / 146 / 21	5.0 / 11.1	3 445 (1 300) / 278 / 88	7.8 / 17.9	1 048 / 177 / 22	10.0 / 28.0
B2a	155 / 69 / 9	5.0 / 14.8	1 005 /243 /82	6.1 / 15.2	822 / 158 / 17	7.6 / 20.2
B2b	112 / 55/ 9	98.0 / 21.0	200 / 73 / 11	7.6 / 20.0	179 / 96 / 13	8.0 / 19.3
B2c	83 / 45 / 8	9.9 / 25.2	69 / 32 / 5	12.6 / 35.3	132 / 74 / 10	10.1 / 33.4
B2d	58 / 31 / 3	15.0 / 49.0	63 / 29 / 4	15.8 / 47.1	78 / 44 / 6	15.0 / 47.5

[a] Columns show mechanism size (N^{rn} / N^{sp} / N^{cd}, number of reactions / species / condensable species) and mean / maximum errors ($\overline{E}$ / E^{max}, %) across the six evaluation scenarios.
[b] Ref: reference mechanism; StP: starting point mechanism; others: reduced mechanisms from different SBR stages. Final mechanisms (A/B with -L, -M, and -H) are highlighted in bold. For C12, early-stage counts are reported as unmerged (merged) for comparison with other mechanisms.
[c] Errors for StP listed here are for the SOA-focused reductions. For multi-target reductions, StP errors are: U8B = 4.7 / 33.3, C12B = 5.2 / 17.3, LIMB = 8.2 / 34.9.

6.1 Size, accuracy, and runtime

Table 6 summarizes the mechanism size and accuracy of representative mechanisms for U8, C12, and LIM obtained during SBR, while Table 7 reports the corresponding computational costs. Overall, SBR produces highly reduced mechanisms under explicitly defined accuracy constraints. Across all VOC classes, reference mechanisms with more than 10^5 reactions are consistently reduced to fewer than 100 reactions (less than 0.1 % of Ref), while mean reduction errors remain below 15 %. Compared to TBR (starting-point for reducing StP mechanism), which reduces mechanisms to 13–20 % of Ref at 5–9 % error, SBR provides about two orders of magnitude further size reduction at comparable error levels.

Figure 11 visualizes the size reduction achieved across PVOCs and reduction configurations. The largest reduction occurs at the earliest SBR stages, followed by more gradual decreases as stricter accuracy constraints are applied, reflecting the staged

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

Table 7. SBR computational cost (T_r in kCPU-h) for the reduction processes, and simulation cost (T_s in CPU-minutes) for running the resulting reduced mechanisms. [a]

Mech.	U8		C12		LIM	
	T_r	T_s	T_r	T_s	T_r	T_s
Ref	–	181.0	–	249.0	–	766.4
StP	0.01	44.4	0.01	51.4	0.02	83.7
SOA-focused reductions (A series)						
A0.1	3.3	1.7	4.3	7.7	17.6	1.5
A1	0.28	1.3	0.84	1.6	0.46	1.2
A2a	2.0	1.2	8.2	1.5	4.4	1.2
A2b	0.59	1.1	2.8	1.3	2.2	1.2
A2c	0.18	1.1	0.44	1.4	2.6	1.1
A3a	0.95	1.1	2.0	1.2	5.1	1.2
A3b	0.46	1.1	0.97	1.2	3.6	1.2
A3c	0.49	1.2	0.77	1.2	2.1	1.2
A3d	0.41	1.1	0.97	1.1	1.6	1.1
Total [b]	8.6	10.0	21.9	16.8	44.9	10.0
Multi-target reductions (B series)						
B0.1	3.94	2.7	4.1	4.3	22.4	2.9
B1	0.62	1.4	0.59	2.9	6.0	1.9
B2a	4.3	1.2	9.9	1.7	28.4	1.6
B2b	1.0	1.1	12.3	1.07	25.8	1.2
B2c	0.72	1.1	1.9	1.1	5.1	1.1
B2d	0.33	0.96	0.54	0.96	4.2	1.1
Total	10.9	7.9	29.4	11.3	91.7	9.2

[a] Computations performed on the NSF NCAR supercomputer using 256 cores for reduction and 64 cores for simulations. T_s reports CPU-minutes aggregated across simulations under 11 scenarios. [b] Total indicates the cumulative cost for SOA-focused or multi-target reductions.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

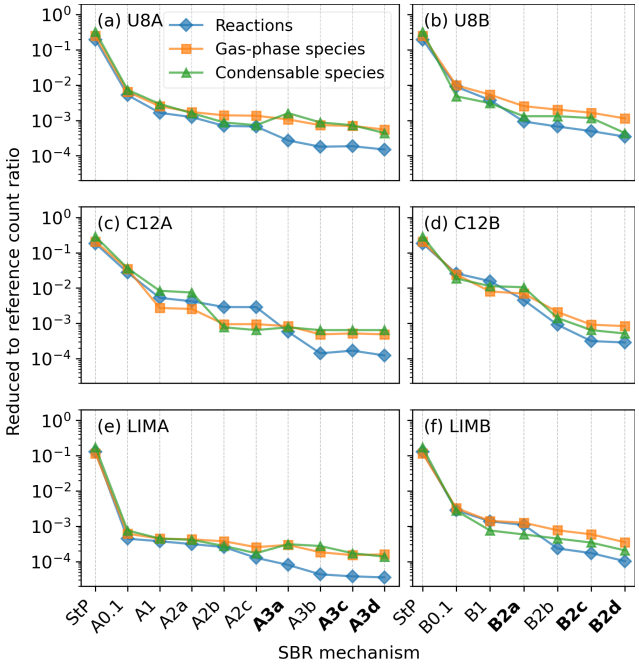

Figure 11. Ratios of reduced to reference mechanism size (by count of reactions, gas-phase and condensable species) for SBR mechanisms obtained under SOA-focused (left panels) and multi-target (right panels) reductions for U8 (top), C12 (middle), and LIM (bottom). Mechanism names shown in bold denote the representative -L, -M, and -H mechanisms.

design of SBR. The reduction scale varies by PVOC, with LIM mechanisms showing the strongest size reduction, followed by C12 and then U8. For each PVOC, SOA-focused reductions achieve greater size reduction than the corresponding multi-target reductions, most clearly in the final mechanisms.

6.1.1 U8 reductions

The U8 reductions illustrate the staged design of SBR and the trade-offs between size and accuracy (example reduction trajectory shown in Fig. S20; Sect. S3.1.2). Both SOA-focused (U8A) and multi-target (U8B) reductions start from StP (32 662 reactions and 6 770 species), already condensed to 20.7 % of Ref with a mean error of 4.7 %. Most of the size decrease occurs during the first SBR cycle (StP to A0.1 or B0.1), where the mechanisms decrease by nearly two orders of magnitude in size while maintaining mean errors below 5 %. This shows that a large fraction of reactions and species can be reduced without immediately compromising accuracy. Stage A2 applies SOA-only reduction, prioritizing the removal of condensable species. Three intermediate mechanisms (A2a, A2b, and A2c) are generated under low, medium, and high error tolerance levels, respectively, and are then further reduced in Stage A3 using matched tolerance sets without the SOA-only restriction. Notably, A2c, generated under the high-tolerance set (15 % mean error and 50 % maximum error), induces mean and maximum errors of only 10.9 % and 16.1 %. As these errors fall within the medium-tolerance range, A2c is also reduced under medium-tolerance

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

settings. As a result, Stage A3 produces mechanisms (A3a to A3d) that represent different balances between size reduction and accuracy. A3b and A3c are both derived under medium tolerances. Although similar in overall size, A3d contains fewer condensable species due to the earlier SOA-only restriction, while A3b contains fewer reactions. This shows that different reduction settings can lead to mechanisms with comparable overall size reductions but with differences in detail. Among these, A3a, A3c, and A3d correspond to the final mechanisms with low, medium, and high error tolerances, respectively. U8SOA-L (A3a: 4.9 % mean error) contains 52 reactions and 29 species (6 condensable species); U8A-M (A3c: 9.1 %) contains 31 reactions and 19 species; and U8A-H (A3d: 14.6 %) contains 27 reactions and 15 species (3 condensable). All three mechanisms are reduced to less than 0.03 % of Ref in size.

Unlike U8A, which focuses on SOA preservation, U8B is designed to retain both SOA and gas-phase complexity. With similar tolerance settings as in U8A, the final U8B mechanisms are larger: U8multi-L (B2a: 155 reactions, 69 species, 5.0 % mean error), U8multi-M (B2c: 83, 45, 9.9 %), and U8multi-H (B2d: 58, 31, 15.0 %). These are approximately twice the size of the corresponding U8A mechanisms, consistent with the additional preservation targets. As the mean error is computed over multiple evaluation targets, SOA errors can be slightly larger in the multi-target reduction than in the SOA-focused reduction under the same tolerance thresholds, resulting in fewer SOA-related species being retained. Consequently, U8B retains a similar number of condensable species as U8A at comparable stages (e.g., A3d and B2d both contain three condensable species). These differences highlight the sensitivity of reduction outcomes to the selected set of evaluation targets.

6.1.2 C12 and LIM reductions

The staged SBRs of C12 and LIM follow the same overall configuration as that of U8, but the efficiency differs across mechanisms. C12 and LIM start from StP containing 18.9 % and 12.0 % of Ref, with mean errors of 9.0 % and 9.4 %, respectively. As these initial errors are higher than for U8, error tolerances for Stage A1 are relaxed to 10 % mean error and 1 % error increase (Table 4). Constrained by the error-increase limit, A0.1 and A1 show a similar error decrease to that in U8A. LIMA reductions are particularly efficient, despite the large size of Ref. The mechanism decreases from 98 794 reactions and 14 417 species to 338 reactions and 77 species after the first cycle (LIMA0.1, 0.048 % of Ref), and to 294 reactions and 56 species after Stage A1 (LIMA1, 0.039 %). These condensed mechanisms are comparable in size to U8A1, indicating that both U8 and LIM contain a limited number of dominant SOA-forming pathways that can be concentrated effectively. In contrast, C12A reductions progress more slowly. After the first cycle, C12A0.1 still retains 4 857 reactions and 1 203 species (2.9 % of Ref), reflecting the more distributed SOA formation in alkane chemistry, which is harder to reduce.

Multi-target reductions display a consistent cross-mechanism pattern. Mechanisms remain comparatively large under strict tolerances but contract sharply once tolerances are moderately relaxed. For C12B, the size decreases gradually through B1 (1.6 %) and B2a (0.46 %) before collapsing at B2b (0.09 %), similar to LIMB, which remains larger at B2a (0.11 %) and then contracts sharply at B2b (0.02 %). In contrast, U8B shows an earlier drop, with the contraction already occurring from B1 (0.38 %) to B2a (0.09 %). This cross-mechanism comparison shows that SBR reduces diverse mechanisms efficiently and progressively, while the pace and extent of reduction depend on the underlying chemical complexity of each mechanism.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

6.1.3 Runtime performance

As shown in Table 7, the reduction procedure is computationally demanding, with total costs ranging from 8.6 to 91.7 kCPU-h. Although the total cost can be substantial, each reduction needs to be performed only once per reference mechanism and configuration, and a single SBR sequence produces multiple reduced mechanisms with different accuracy levels. These can subsequently be applied for different modeling purposes without additional reduction cost. SOA-focused reductions are consistently faster than multi-target reductions, even though the former include more stages (A1–A3 versus B1–B2). For large and chemically complex mechanisms, multi-target reductions incur substantially higher computational cost, reaching approximately twice that of SOA-focused reductions for LIM (44.9 kCPU-h for LIMA compared to 91.7 kCPU-h for LIMB). Although the per-evaluation cost difference between evaluating one or multiple evaluation targets is small, tracking additional evaluation targets introduces more evaluation criteria, leading to more frequent rejection of SBR candidates and thereby slowing the reduction. As a result, the reduction cost is more sensitive to evaluation target selection than to the number of stages in the presented reductions.

The distribution of computational cost across stages also varies. The first cycle of Stages A1 and B1 dominates the stage cost, accounting for 79.0 % (LIMA0.1) to 97.5 % (LIMB0.1). This reflects the large mechanism size at the beginning, when identification and evaluation of reduction candidates are most expensive. For mechanisms where substantial size reduction occurs in early stages (U8A, U8B, and LIMA), later stages (Stages A2–A3 and B2) contribute only about 55 % of the total cost. In contrast, later stages remain costly for larger mechanisms (C12A, C12B, and LIMB), contributing nearly 80 % of the total cost, particularly when strict error tolerances are applied.

Overall, the computational cost depends strongly on both the mechanism and the reduction settings. Key factors include mechanism size, whether the evaluation focuses solely on SOA or includes additional metrics, the stringency of early-stage error tolerances, and the rate at which tolerances are relaxed. Larger mechanisms therefore require more computational resources, particularly when broader evaluation targets and stricter error tolerances are imposed. The computational benefit of reduction is evident in simulation runtime. Reference simulations across 11 scenarios require 181–766 CPU-min, scaling with mechanism size. After early-stage reductions (A1 and B1), simulation costs decrease to about 1 CPU-min for both SOA-focused and multi-target mechanisms, corresponding to speed-ups exceeding two orders of magnitude.

6.2 Preservation of key metrics

Sect. 6.1 evaluates SBR performance using reduction errors across six SBR evaluation scenarios, whereas the following analysis uses all 11 evaluation scenarios, consistent with TBR performance evaluation (Sect. 5). SBR accuracy is assessed using mean NMB for SOA-related metrics (Sect. 6.2.1), as well as for gas-phase metrics with additional assessment of tracer species (Sect. 6.2.2). NMB distributions for the U8 reduced mechanisms are shown in Fig. 12, with corresponding results for C12 and LIM in Figs. S21–S22 (Sect. S3.2).

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

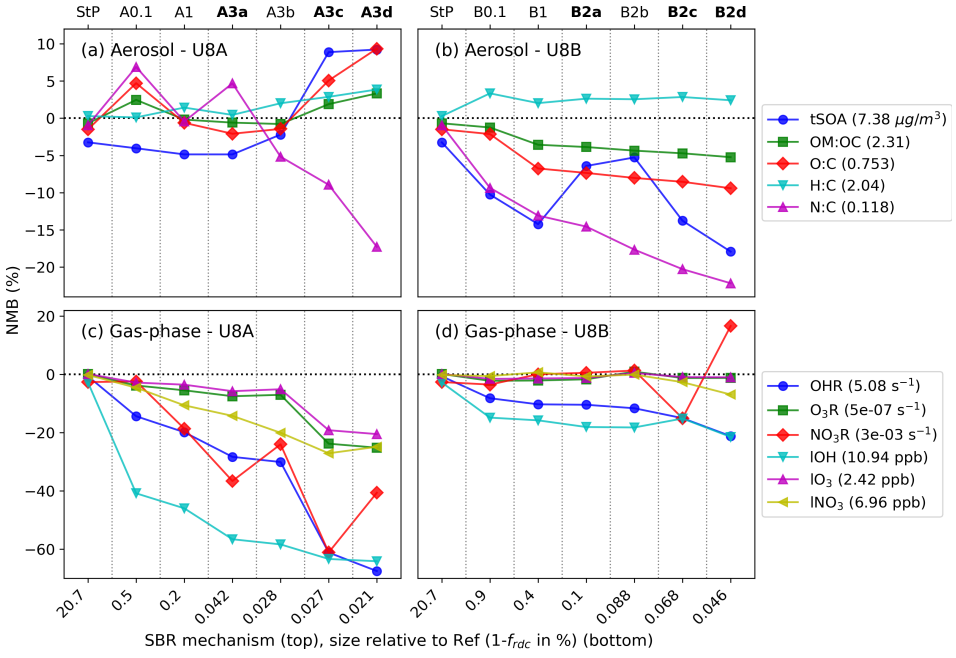

Figure 12. Reduction accuracy (mean NMB) of key metrics simulated with reduced mechanisms obtained from U8A (left panels) and U8B (right panels) reductions across 11 scenarios, together with their relative size fraction compared to the reference mechanism (Ref). Shown metrics include SOA mass (tSOA) and ratios (tSOA, OM:OC, O:C, H:C, and N:C) in the top panels, and gas-phase reactivities (OHR, O_3R, NO_3R) and related tracer species (lOH, lO_3, lNO_3) in the bottom panels. Reference mean values are indicated in the legend. Mechanism names shown in bold denote the representative -L, -M, and -H mechanisms.

6.2.1 SOA-related metrics

SOA-related chemistry preservation is evaluated using tSOA and elemental SOA ratios (e.g., shown for U8 in Figs. 12 (a)–(b)). Among the reduced mechanisms, StP reduced by TBR exhibits the lowest biases. As SBR proceeds with increasingly aggressive size reduction and higher error tolerance settings, deviations increase but remain acceptable when evaluated relative to the achieved size reduction. Across all VOC classes and SBR configurations, tSOA and biases in OM:OC, O:C, and H:C remain limited. H:C and OM:OC exhibit the narrowest bias ranges, generally within ±5 %. tSOA and O:C show larger variability (typically within ±10–20 % NMB), with the largest deviations occurring for the most reduced mechanisms (SOA-focused: U8A3d at -9.2 %; multi-target: U8B3d at 17.9 %). These results demonstrate that tSOA remains an effective evaluation target for preserving bulk SOA formation under SBR. N:C, however, exhibits the highest sensitivity to SBR, with relatively large biases observed even at early stages (e.g., U8A0.1 with N:C NMB at 6.9 %), reflecting reduced preservation of nitrogen-containing SOA components. As SBR evaluation in all reductions presented includes only tSOA as a bulk constraint, these results suggest that this constraint alone may be insufficient to preserve detailed SOA composition. Additional SOA-related evaluation targets or constraints would therefore be required for SBR to better control compositional preservation.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

Systematic responses to SBR evaluation criteria are also evident. Under comparable reduction accuracy control, SOA-focused reductions induce smaller biases than multi-target reductions. These differences arise from the SBR candidate evaluation, in which reduction candidates are assessed using reduction errors computed across multiple evaluation targets (Eq. C6), allowing partial compensation among metrics while remaining bounded by prescribed maximum error tolerances. As the SBR error metric is based on absolute errors rather than signed bias, the resulting reduced mechanisms exhibit shifts in bias and a mixture of positive and negative NMB values, reflecting adjustments among competing chemical pathways under constrained optimization. Meanwhile, bias patterns vary with mechanism complexity. U8 and LIM exhibit better SOA preservation than C12, reflecting the increased difficulty of preserving more distributed and individually weaker SOA-forming pathways in C12, consistent with the observed size–accuracy relationships in Table 6.

6.2.2 Gas-phase metrics and tracers

Gas-phase chemistry preservation is evaluated using NMB for bulk gas-phase reactivities and associated tracer species. The preservation of these quantities varies strongly between SOA-focused and multi-target reductions, highlighting the importance and effectiveness of SBR evaluation target selection. Across all PVOCs, multi-target reductions consistently produce smaller and more centered biases than SOA-focused reductions, with the contrast becoming most pronounced for highly reduced mechanisms. For U8 shown in Figs. 12 (c)–(d), the most reduced multi-target mechanism (U8B2d) exhibits an OHR NMB of (-21.2 %), compared to (-67.5 %) for the corresponding SOA-focused mechanism (U8A3d). These bias distributions demonstrate that the inclusion of gas-phase tracers in multi-target reductions substantially improves the constraint of gas-phase chemistry compared to SOA-focused reductions.

This systematic improvement in gas-phase metric preservation across PVOCs reflects the role of gas-phase tracers in constraining dominant oxidant budgets and reaction pathways. Although gas-phase reactivities are not directly included as evaluation targets, the use of loss tracers that track oxidant consumption (e.g., lOH for OHR, lO_3 for O_3R, and lNO_3 for NO_3R) limits deviations in major oxidant sinks within the reduced mechanism, thereby indirectly constraining the associated reactivities and improving gas-phase chemistry preservation. The effectiveness of tracer-based constraints varies by tracer and VOC class, reflecting the non-linear relationship between tracer species. For U8, NO_3R NMB changes by more than 30%, from -15.1 % in U8B2c to 16.6 % in U8B2d, while lNO_3 NMBs change by less than 5 % (from -2.7 % to -6.9 %), suggesting weak correlation between NO_3R and lNO_3 in this case. In contrast, O_3R and lO_3 exhibit closely aligned NMB distributions (e.g., -25.3 % and -20.5 % for U8A3d; -1.3 % and -1.0 % for U8B2d). These two examples show that tracer–reactivity correspondence varies by tracer and underlying gas-phase chemistry.

Figures 13(a)–(f) show NMB distributions for NO_2 and HO_2 tracers across PVOCs (lNO_2, pNO_2, lHO_2, and pHO_2), along with NMBs for the corresponding mixing ratios. Across all PVOCs, tracer NMB is significantly lower for multi-target reductions than for SOA-focused reductions, as a direct result of including these tracers in the SBR evaluation target sets. Consistent with the tracer–reactivity distributions discussed above, this improved tracer preservation does not translate uniformly to the corresponding species mixing ratios. NO_2 biases remain small even in SOA-focused reductions (up to 15.6 % for LIMA and 29.5 % for LIMB), whereas HO_2 shows clearer improvement when explicitly constrained (e.g., NMB changing from -62.3 %

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

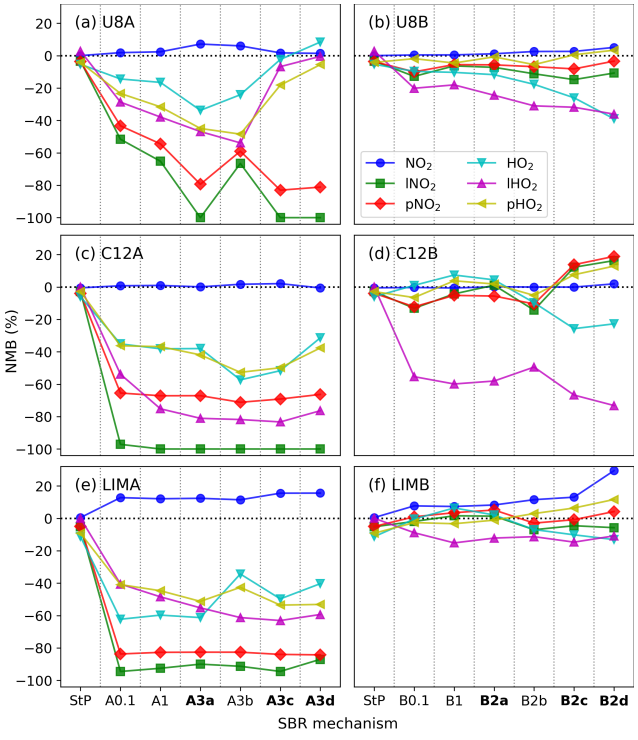

Figure 13. Reduction accuracy (NMB) of NO_2 and HO_2 and their corresponding tracer species (lNO_2, pNO_2, lHO_2, and pHO_2) used in the SBR multi-target reduction. Results are shown for SOA-focused (left panels) and multi-target (right panels) reduced mechanisms for U8 (top), C12 (middle), and LIM (bottom) across 11 scenarios.

for LIMA to -13.1 % for LIMB). The results are consistent with the chemical coupling in which NO_2 is controlled primarily by NO_x–O_3 chemistry, with relatively minor contributions from VOC degradation, whereas HO_2 is produced largely within VOC oxidation pathways. Accordingly, indirect constraints based on tracer inclusion are effective for multi-target reduction, but reliable tracer representation depends on selecting species that are chemically well coupled to the targeted processes.

When multiple evaluation targets are applied, preservation of related metrics is partly achieved through compensation among targets rather than strict pathway- or target-level control, reflecting an inherent but configurable trade-off in the current SBR evaluation configuration. Consequently, localized bias amplification or offset may occur (e.g., a large mean lHO_2 bias of -51.8 % across all C12B mechanisms). Such effects can be mitigated through adjustments to the evaluation configuration, particularly by adopting stricter maximum error tolerances and improving the selection of evaluation targets. Nevertheless, multi-target reductions consistently outperform SOA-focused reductions in preserving gas-phase chemistry while maintaining SOA formation within acceptable bounds, demonstrating that tracer-based multi-target SBR provides an effective and flexible framework, with potential for further refinement of tracer selection.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

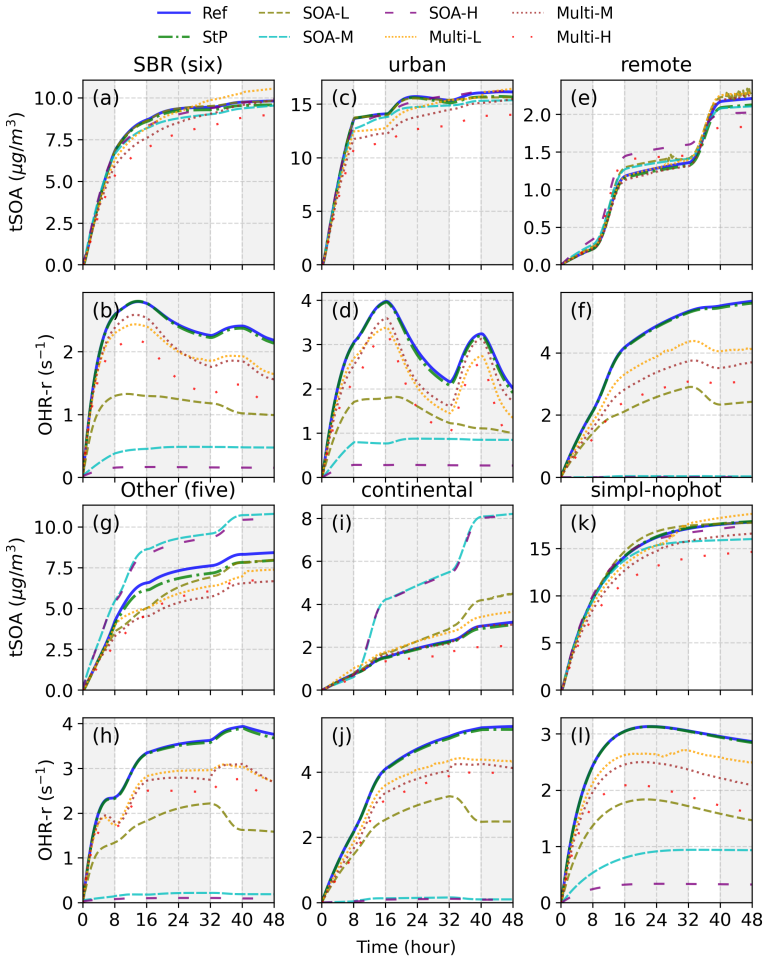

Figure 14. Time evolution of tSOA (top and third rows) and OHR-r (second and bottom rows) simulated with U8 mechanisms. Results are shown as mean across SBR evaluation scenarios (SBR(six); panels a–b) and non-SBR scenarios (Other(five); panels g–h), together with representative scenario-specific simulations (urban and remote from SBR; continental and simpl-nophot from Other; panels c–f and i–l). Mechanisms include Ref, StP, and SOA-focused and multi-target SBR reductions at three error tolerance levels (L, M, and H). Gray shading denotes time periods (0–8, 16–32, and 32–40 h), corresponding to nighttime (little or no solar radiation) in scenarios with photolysis based on solar position.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

6.3 Time variation preservation

Temporal evolution is analyzed to evaluate whether SBR preserves time-dependent chemical behavior beyond aggregated metrics. Two key diagnostics are examined: tSOA representing SOA formation, and OHR-r representing gas-phase chemical activity from oxidation products. Both metrics are accumulated over reducible products and therefore directly reflect the impact of mechanism reduction on chemically active pathways. Figure 14 presents U8 time series, including SOA-focused and multi-target reductions at three error tolerance levels. Results are organized by the six scenarios used during SBR candidate evaluation (SBR(six)) and the five remaining scenarios (Other (five)). For each set, both scenario-mean time series and two representative individual scenarios are shown. Corresponding results for C12 and LIM, as well as for the remaining U8 scenarios, are provided in Figs. S23–S25 (Sect. S3.2.1), with additional chemistry-driven interpretation for U8.

Differences in the magnitude and temporal structure of tSOA and OHR-r across scenarios reflect contrasts in oxidant availability, NO_x levels, and photochemical treatment. Except for scenarios with prescribed high photolysis (urban-highphot and simpl-highphot), simulations generally begin at midnight under dark conditions, with the initial 0–8 h dominated by nighttime chemistry driven by O_3 and NO_3 oxidation. For scenarios with diurnally varying photolysis, high-NO_x scenarios generally show more rapid oxidation and larger tSOA formation at the beginning of the simulations (e.g., urban, Fig. 14(c)), while low-NO_x scenarios lead to weaker nighttime chemistry and stronger dependence on daytime photochemical pathways (e.g., remote, Fig. 14(e)). Fixed photolysis scenarios exhibit smoother temporal evolution due to constant oxidation conditions (e.g., simpl-nophot, Fig. 14(k)).

6.3.1 Performance under SBR and non-SBR scenarios

Despite the diversity of chemical regimes, temporal features are largely retained in the reduced mechanisms. For the six SBR evaluation scenarios, all reduced mechanisms across PVOCs reproduce the time evolution of tSOA closely relative to Ref, with only minor deviations even for the most reduced reductions (SOA-H and multi-H). This response is consistent with previously discussed reduction errors and NMB distributions, as tSOA is explicitly included as an SBR evaluation target in all reductions. Preservation of OHR-r shows stronger dependence on reduction configuration. Multi-target reductions generally capture both the magnitude and diurnal structure of OHR-r, although increasing error tolerance leads to systematic underestimation. SOA-focused reductions, however, exhibit a rapid collapse of OHR-r after the initial oxidation period, reflecting the absence of gas-phase constraints. This behavior highlights the importance of including gas-phase metrics, such as reactivities or related tracers, when preservation of gas-phase chemistry is desired.

For scenarios not included in SBR candidate evaluation, preservation remains within reasonable bounds for tSOA across all reductions and for OHR-r in multi-target mechanisms, but performance becomes more scenario-dependent. Scenarios chemically similar to those used for training remain well reproduced, whereas scenarios with distinct oxidant or photochemical regimes show larger deviations. Preservation of OHR-r varies among PVOCs, reflecting differences in dominant oxidation pathways. For U8, the simpl-nophot scenario shows close agreement with SBR scenario results for both tSOA and OHR-r across all reductions (Fig. 14(k)–(l)). In contrast, under the continental scenario, SOA-focused reductions at medium and

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

high error tolerances (SOA-M and SOA-H) show substantial divergence in tSOA following sunrise (Fig. 14(i)–(j)), indicating sensitivity to daytime photochemical SOA formation pathways that are not constrained in those mechanisms.

6.4 Reaction pathway visualization

Reaction pathway visualizations provide direct mechanistic inspection of how SBR reshapes oxidation chemistry while preserving predefined evaluation targets. The examples presented demonstrate that pathway simplification remains chemically consistent across reductions, with differences primarily attributable to the selected error tolerance and evaluation target configuration. To examine these effects, reaction pathways for U8 SOA-focused reductions at high and medium error tolerances (U8SOA-H and U8SOA-M; Figs. 15(a)–(b)) are discussed to assess the impact of accuracy constraints on pathway retention. Comparison with the corresponding multi-target reduction (U8Multi-H; Fig. 16) is used to assess the influence of evaluation target selection. Mappings between original species and new surrogate species resulting from lumping are provided in Tables S6–S8 (Sect. S3.4) for the pathways shown (including species reduced by replacement to lumped surrogates, when applicable). Surrogate species are prefixed with "m" followed by the dominant parent species name, and their structures are used in the reaction pathways. Additional reaction pathways generated with the same tolerance settings for the C12 and LIM reductions are provided in the Supplement (Sect. S3.3; Figs. S26–S28 for C12 and Figs. S29–S31 for LIM), providing additional examples across PVOCs with differing chemical complexity.

6.4.1 SOA-focused preservation

The U8SOA-H mechanism (Fig. 15(a)) exemplifies the most aggressively reduced mechanism for SOA formation from 1-octene (U08000). Only three condensable products remain in the reduced mechanism, compared with 6 836 in the reference mechanism. These include two organic nitrates (GH6002 and mHN800n) associated with O_3- and NO_3-initiated degradation, and one hydroperoxide (mHO8002) formed through OH-initiated pathways. Despite the strong reduction, key reaction types are preserved in the U8SOA-H mechanism. Following initial oxidation of U08000, competing fates of RO_2 radicals are retained through reactions with NO, HO_2, and RO_2, maintaining essential branching in the oxidation chemistry. For example, the RO_2 surrogate m2O8001 participates in multiple channels, forming alkoxy radicals (1O8001), secondary RO_2 species (m2O6000), and closed-shell organic nitrates (mNO8001), while also contributing to the formation of small carbonyls such as formaldehyde (CH_2O). Subsequent reactions originating from these pathways lead to the formation of all condensable products. Terminal non-condensable products that do not form further reducible species are also retained (e.g., mHO8002), maintaining chemical closure and mass balance without introducing additional SOA-relevant complexity. Together, these preserved pathways demonstrate that, even with a limited number of surrogate species, the reduced mechanism maintains key gas-phase processes influencing SOA formation.

In contrast, gas-phase degradation is not intended to be accurately represented in U8SOA-H. O_3-initiated chemistry is limited to pathways that remain necessary under the tSOA constraint to balance CH_2O and CH_2OO concentrations, consistent with previous discussion that SOA-focused reductions can still partially reproduce gas-phase chemistry. As a result, a limited subset of gas-phase pathways is preserved to support tSOA accuracy, indirectly influencing oxidant and inorganic budgets. Overall,

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

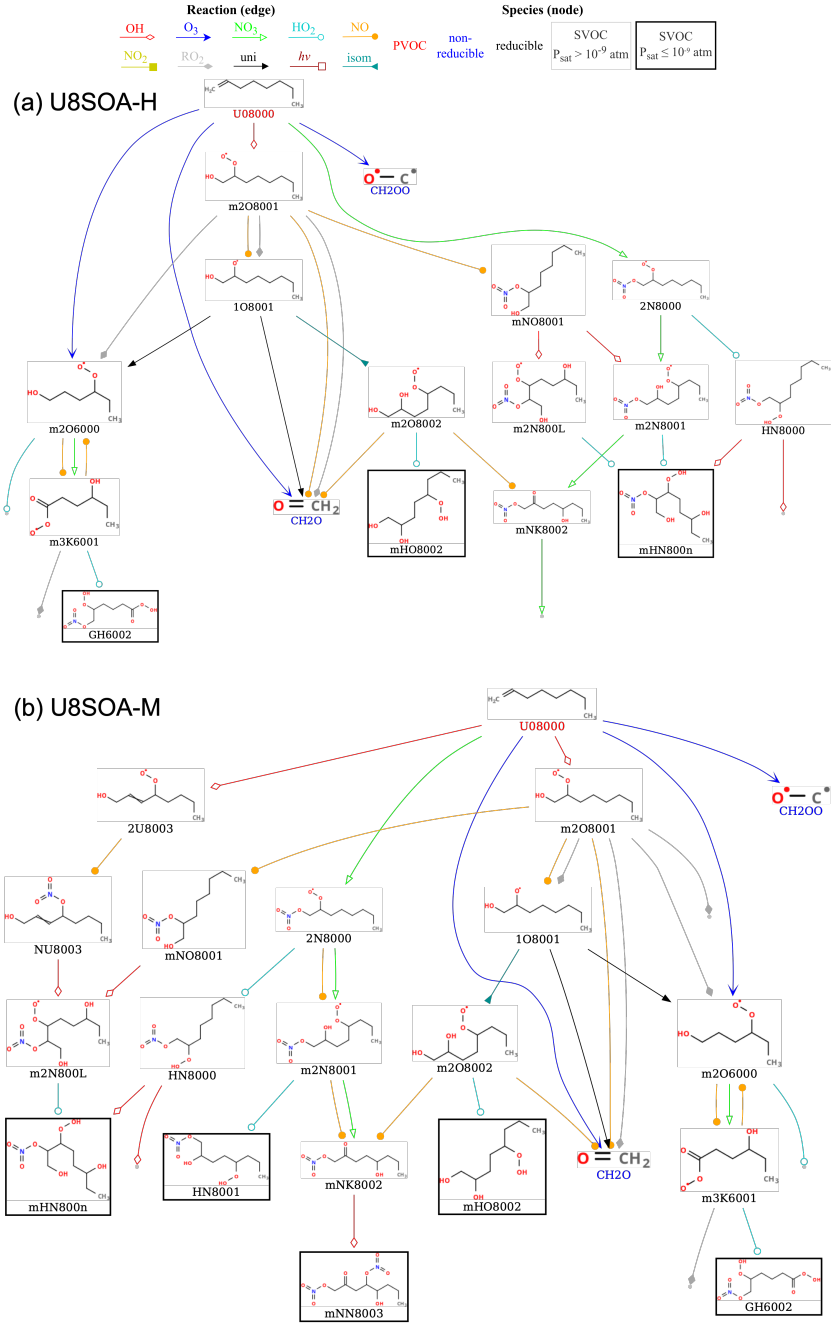

Figure 15. Reaction pathways of the U8 SOA-focused reduced mechanisms with high and medium error tolerances (U8SOA-H or U8A3d, panel a; U8SOA-M or U8A3c, panel b). Species are nodes with names (six-letter GECKO-A name from Ref.; lumped surrogates prefixed with "m") and structures. PVOCs are red, non-reducible species blue, and reducible species black; reducible SVOCs are boxed (darker = lower volatility). Reactions are colored edges by type: OH (red diamond), O_3 (blue V), NO_3 (green triangle), HO_2 (cyan circle), NO (orange circle), RO_2 (gray diamond), unimolecular ("uni", black tee), photolysis ("*hv*", brown square), and isomerization ("isom", darkcyan inverted triangle). Terminal products are connected to gray dots to indicate further degradation.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

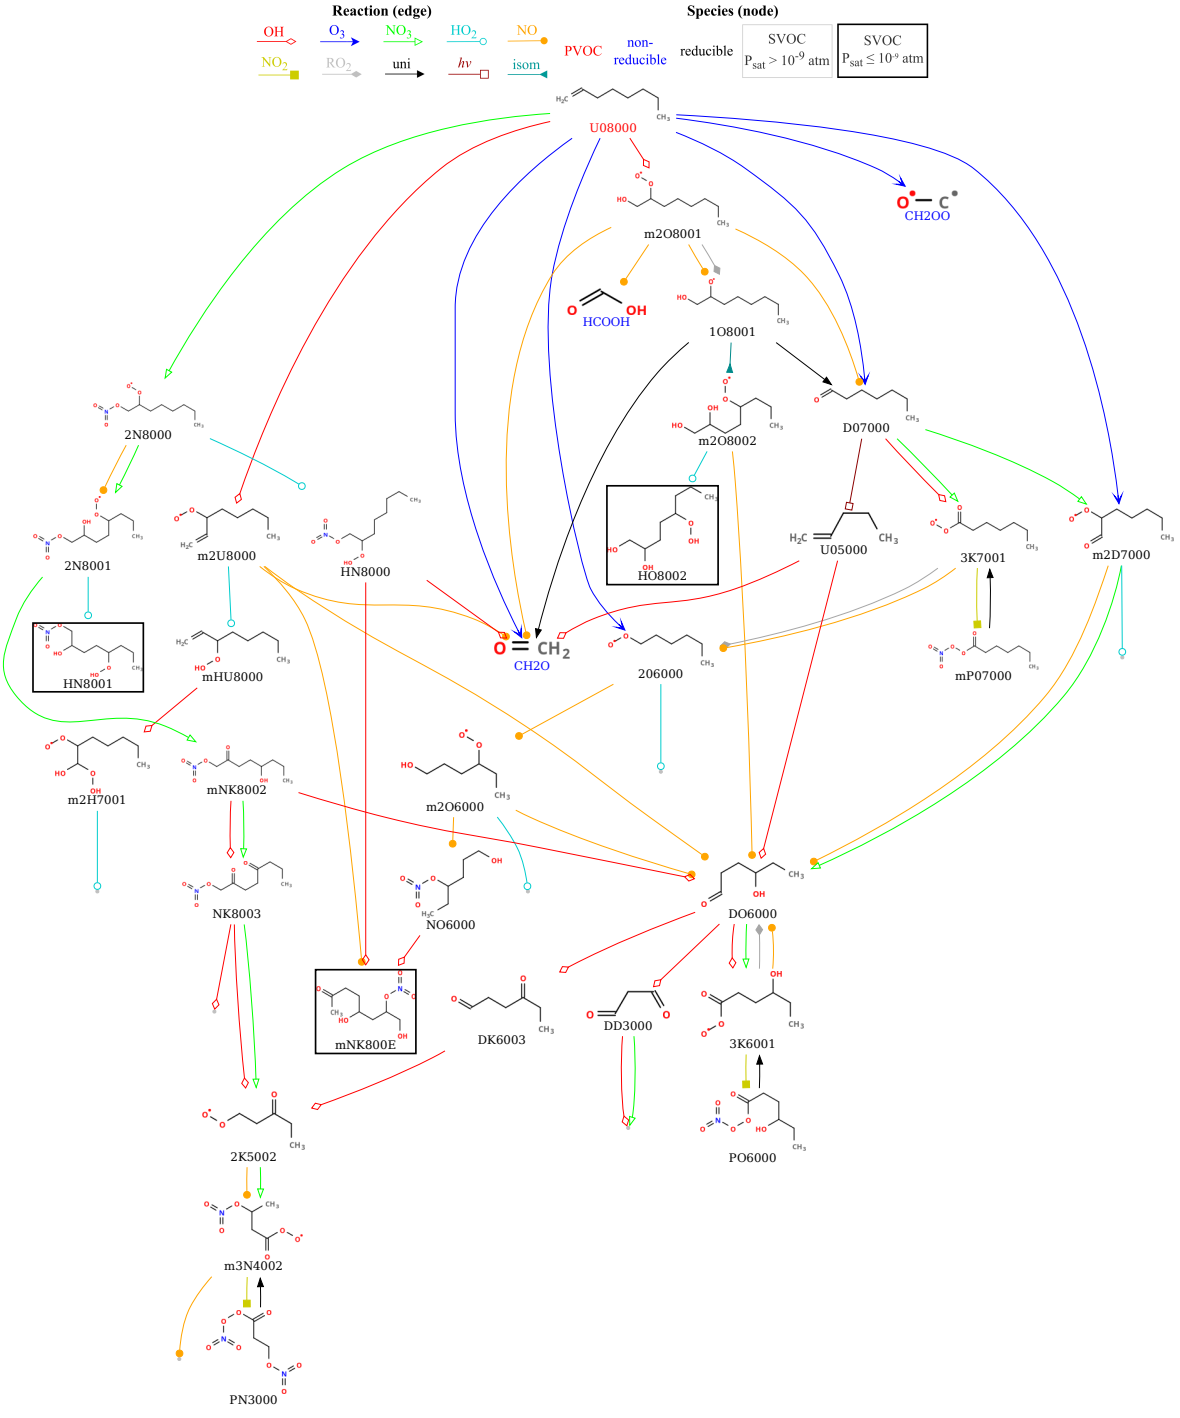

Figure 16. Reaction pathways of the U8 multi-target reduced mechanism with high error tolerances (U8Multi-H or U8B2d). Graph node and edge conventions are the same as in Fig. 15.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

the mechanism is limited to a maximum node depth of four, with minimal branching and direct connections to low-volatility products.

6.4.2 Influence of error tolerances

The differences between U8SOA-M (Fig. 15–(b)) and U8SOA-H illustrate the impact of increasing error tolerance in the late-stage reduction. Both mechanisms follow the same reduction pathway through Stage A2, after which they diverge in Stage A3 according to the applied tolerance (30 %/10 % and 50 %/15 % for mean and maximum tSOA errors across the SBR evaluation scenarios, respectively). Despite these differences, U8SOA-M and U8SOA-H retain highly similar overall pathway structures. Most pathways and species in U8SOA-M are preserved in U8SOA-H, and all lumping surrogate species are identical, since both mechanisms derive from the same Stage A2 reduction.

Only a few pathway-level differences are introduced in U8SOA-H relative to U8SOA-M. These include the further reduction of condensable species mNN8003 via species removal and HN8001 via jumping (m2N8001 directly to mHN800n) in Stage A3; the reduction of the U08000 OH-initiated pathway forming 2U8003, where 2U8003 is first jumped to NU8003 and the remaining NU8003 reactions are removed; and the removal of one RO_2 reaction of m2O8001. These changes demonstrate how GENOA v3 systematically trims or merges weakly influential branches as error tolerance increases. At late-stage SBR, the mechanism is already compact and further reductions are strongly mechanism dependent. When additional reductions would induce errors exceeding the allowed tolerance, different tolerance settings can produce near-identical reduced mechanisms. Although the exact reduction direction is determined by SBR evaluation criteria, the resulting mechanisms remain structurally consistent when error growth limits further reduction.

6.4.3 Influence of evaluation targets

The differences between U8SOA-H and U8multi-H (Fig. 16) demonstrate the influence of evaluation target selection under identical error tolerances and SBR evaluation scenarios. Although the two reductions diverge from the beginning of SBR due to different targets, both mechanisms preserve comparable SOA-forming pathways, as tSOA is included in both evaluation target sets. The dominant first-generation RO_2 species formed during octene oxidation that contribute to SOA formation are largely consistent between the two mechanisms. In both reductions, the OH-initiated product m2O8001 and the NO_3-initiated product 2N8000 are retained; for O_3-initiated oxidation, the major product 2O6000 is retained explicitly in U8multi-H and represented by a lumped surrogate (m2O6000) in U8SOA-H. These SOA-relevant first-generation pathways subsequently form condensable products that correspond closely between the two mechanisms, with HO8002, HN8001, and mNK800E in U8multi-H corresponding, respectively, to mHO8002, mHN800n, and GH6002 in U8SOA-H, and exhibiting similar functional group composition and volatility.

Beyond SOA formation, inclusion of gas-phase tracer species in U8multi-H leads to broader pathway retention. U8multi-H preserves substantially more intermediates and reactions associated with oxidant budgets and NO_x chemistry, nearly doubling the mechanism size and extending maximum node depth to approximately eight, compared to U8SOA-H. Additional first-generation products are retained in U8multi-H, including m2U8000 from OH-initiated oxidation, as well as D07000 and

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

m2D7000 from O_3-initiated reactions, which further form multiple nitrate- and peroxy-containing products that remain explicitly represented. These pathways influence tracer species such as lNO_3, pNO_2, and lHO_2, but contribute little to SOA mass. The multi-target configuration also preserves peroxy acyl nitrate (PAN)-related terminal products (e.g., PN3000 and P06000), capturing NO_2 buffering and release processes that are absent from the SOA-focused reduction. 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.

7 Reduction methodology, limitations, and future directions

7.1 Comparison of TBR and SBR

After discussing the reduction performances of TBR and SBR approaches separately, their methodological differences and outputs are compared here. Both follow a candidate-based structure but differ in scope, accuracy control, and computational demand. TBR applies rapid, threshold-driven reductions across the entire mechanism, typically by removing species and reactions. It is computationally inexpensive and particularly valuable for large mechanisms, where it can quickly remove less influential components. Its performance depends on the choice of strategies, threshold values, and their application order. Although this process does not include built-in evaluation, optional assessment with metrics similar to those in Sect. 3.3.3 can help balance trade-offs between size, complexity, and accuracy. SBR instead relies on iterative, simulation-informed evaluation of candidates using quantitative metrics. This approach is more computationally demanding but provides finer accuracy control and results in deeper reductions. It is most effective for small or moderate-sized mechanisms, or when TBR alone cannot achieve the desired compression. Key parameters include those used for candidate identification, such as species grouping and ordering, and those used for evaluation, including the choice of evaluation scenarios and metrics.

The presented examples illustrate these methodological differences. For U8, TBR (U8StP) retained 20.7 % of the original size, while the SBR decreased it further to below 0.05 % (U8SOA-H and U8multi-H). Such reductions cannot be achieved by TBR alone while preserving reasonable accuracy. However, 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 user-defined objectives.

7.2 Limitations and sources of uncertainty

Although GENOA v3 effectively reduces highly detailed GECKO-A mechanisms while preserving key evaluation targets, the resulting reduced mechanisms remain subject to several limitations that should be considered in their application. One inherent source of uncertainty is the accuracy and structural assumptions of the reference mechanism itself, which are directly propagated into the reduced mechanisms. Beyond this dependency, reduction outcomes are primarily influenced by scenario selection and methodological simplifications within the reduction workflow.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

7.2.1 Scenario dependence and chemical representation

Reductions in GENOA v3 are inherently scenario dependent. Predefined evaluation scenarios are used throughout the workflow, including during the estimation of environment-sensitive properties and the evaluation steps during both TBR and SBR. 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. Increasing the number of scenarios can improve coverage of atmospheric variability, but directly raises computational cost through longer candidate evaluation during SBR. A practical approach is therefore to begin with small, condition-specific scenario sets targeting well-defined chemical regimes and expand them iteratively as broader applicability is required.

Additional uncertainty arises from simplified chemical representations adopted during reduction. In the current implementation, reference concentrations are represented using fixed mean values across scenarios, even though scenario-specific concentrations can differ by several orders of magnitude. Scenario-dependent formulations or concentration ranges may provide a more realistic representation but require further evaluation to balance improved fidelity against computational expense. As another example, although some TBR strategies are chemically informed, they are currently applied uniformly across reaction or species types (e.g., using the same node-depth limit for all species). These strategies could be made more targeted and effective by incorporating more detailed chemical rules tailored to the input mechanism (e.g., applying different thresholds to closed-shell and radical species).

7.2.2 Configuration sensitivity and carbon conservation

The reduction outcomes are also sensitive to configuration choices. For TBR, this primarily involves the choice of reduction strategies and associated thresholds. For SBR, both candidate identification and evaluation influence the reduction. SBR Candidate identification can be adjusted by modifying existing strategies (e.g., alternative lumping or surrogate construction approaches), while candidate evaluation and selection have a strong influence on reduction outcomes and should be guided by the intended application (e.g., through appropriate choices of evaluation targets, tolerances, and scenarios). In the current implementation, the reduction workflow 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 (Sect. S3.5). While such carbon loss is acceptable for target-focused applications, future developments could explore placing greater emphasis on lumping-based strategies with explicit carbon control.

7.2.3 Computational cost and scalability of SBR

The computational demand for SBR can increase rapidly for very large mechanisms (e.g., Table 7), as candidate evaluation relies on repeated box-model simulations under predefined scenarios. The cost therefore scales with mechanism size, number

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

of evaluation scenarios, and accuracy constraints, which may limit direct application to extremely detailed mechanisms. More aggressive TBR prior to SBR, or the introduction of an intermediate reduction stage, may help address this limitation. Such an intermediate approach would follow a workflow similar to SBR but would rely on yield- or budget-based estimates rather than full simulations to assess preservation of evaluation targets. Such estimates are computationally inexpensive and can provide reasonable accuracy for early-stage candidate evaluation, although they do not capture nonlinear interactions and time-dependent pathway competition that can be explicitly resolved in box-model simulations. This intermediate step could therefore serve as a practical bridge between TBR and SBR, enabling efficient reduction of large mechanisms before applying SBR at later stages to smaller mechanisms where more detailed chemical interactions can be resolved.

7.3 Community use and future directions

Despite these limitations, GENOA v3 is designed to support flexible, user-configurable reduction workflows, allowing strategies, scenarios, and evaluation settings to be adapted to specific applications and modeling needs. Beyond individual case studies, one potential direction is the development of predefined sets of reduced mechanisms optimized for representative environments, such as urban atmospheres or wildfire plumes. Such collections could function as ready-to-use mechanism libraries for regional or global models and lower the barrier to adopting reduced chemistry in large-scale simulations. Future methodological improvements may focus on refining chemically informed reduction strategies, strengthening candidate selection, and increasing evaluation efficiency to achieve optimal reductions with lower computational cost while preserving the metric-driven framework. Advances in scenario design and evaluation methods, with possible support from ML, may further improve reduction performance.

Community involvement is expected to play an important role in these developments. By enabling user-defined configurations and making reduction settings transparent and traceable, GENOA v3 provides a platform for testing new strategies, sharing reduced mechanisms, and incorporating feedback from diverse modeling contexts. Continued development will emphasize extensibility, transparency, and reproducibility, allowing the framework to evolve alongside community needs and advances in atmospheric chemistry modeling.

7.4 Integration with machine learning

Conceptual similarities can be found between ML approaches for mechanism reduction and GENOA v3 reduction approaches in their shared objective of reducing system complexity while preserving key behaviors under explicit accuracy constraints. Some reduction strategies used in GENOA v3 are also graph-based, an approach that is commonly adopted in recent ML studies. However, GENOA reductions retain detailed chemical representation and mechanistic interpretability, producing a series of reduced mechanisms that are smaller in size yet chemically consistent, depending on the reduction settings applied. These features are often not preserved in current ML-based approaches. The complementary strengths of the two approaches suggest potential opportunities for integration. GENOA-reduced mechanisms provide chemically consistent reduced mechanisms that could serve as training datasets for ML models. Conversely, ML techniques may help accelerate candidate identification and evaluation or offer more efficient alternative strategies. Hybrid frameworks could therefore combine the interpretability of

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

GENOA with the efficiency of ML, offering a pathway toward scalable and reliable mechanism reduction in three-dimensional atmospheric models.

8 Conclusions

This paper introduces the GENOA v3 framework designed for optimized atmospheric chemical mechanism reduction. As the third iteration in the family of GENOA reduction algorithms, it introduces a significant evolution, enabling the systematic top-down reduction of highly detailed VOC mechanisms generated by GECKO-A, which can encompass up to millions of species and reactions.

Two reduction approaches are implemented in GENOA v3: threshold-based reduction (TBR), which applies strategy-driven pruning, and simulation-based reduction (SBR), which performs iterative, graph-aware, accuracy-controlled refinement coupled with box-model evaluation. TBR provides rapid and efficient initial pruning of large mechanisms at low computational cost, while SBR achieves more extreme reduction at higher computational demand, with essential complexity preserved. Together, this dual workflow enables efficient mechanism reduction while maintaining accuracy for both SOA formation and gas-phase chemistry.

Applied to 29 GECKO-A degradation schemes of diverse precursor classes, including n-alkanes (C5–C16), 1-alkenes (C5–C16), and monoterpenes (α-pinene, β-pinene, limonene, camphene, carene), TBR achieved size decreases of 23–65 % (46 % mean) with conservative strategies (T6) and 68–96 % (80 % mean) with more aggressive strategies (T10). Accuracy preservation depended on strategy aggressiveness: tSOA was well preserved under T6 with median and mean errors of 2.1 % and 0.73 %, whereas T10 increased these to 12.2 % and 3.7 %, driven by a small number of high-error reductions. A similar response was observed with preservation of gas-phase chemistry targets, with aggregated metric-level mean errors of 1.9–7.5 % under T6 compared to 2.7–25.9 % under T10, reflecting increasing sensitivity to pathway reduction, largely driven by TBR removal strategies. Reduction efficiency generally increased with precursor carbon number and mechanism size, whereas small mechanisms (C5–C6 and U5–U6) and structurally complex mechanisms such as monoterpenes were more sensitive to fixed thresholds, leading to larger deviations when the additional environmentally sensitive strategies in T10 were applied.

SBR further reduced mechanisms through staged candidate identification and evaluation under explicit accuracy constraints. Reference mechanisms exceeding 10^5 reactions were consistently reduced to fewer than 100 reactions (< 0.1 % of the reference size) while maintaining mean errors below 15 %. The largest reduction was obtained for limonene, reduced from 762 747 reactions and 125 314 species to 27 reactions and 20 species (0.006 %). Reduction performance depended strongly on evaluation target selection. SOA-focused reductions produced the most reduced mechanisms by constraining SOA-related metrics, whereas multi-target reductions retained greater chemical complexity and substantially improved preservation of oxidant budgets, typically requiring mechanisms about twice as large. For U8, a highly reduced SOA-focused configuration exhibited an OHR bias of -67.5 %, while a larger multi-target configuration reduced this bias to -21.2 %, demonstrating that tracer-based constraints effectively limit deviations in major oxidant sinks.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

Together, these results show that TBR provides efficient first-stage pruning with predictable accuracy degradation as strategies become more aggressive, whereas SBR enables accuracy-controlled refinement aligned with user-defined evaluation targets. The combined workflow preserves dominant chemical pathways while allowing explicit trade-offs between mechanism size and chemical accuracy, providing a configurable and chemically interpretable framework for generating compact mechanisms suitable for applications ranging from box models to three-dimensional atmospheric simulations.

Code and data availability. GENOA v3.0 is publicly available on GitHub (https://github.com/tool-genoa/genoa), with documentation hosted via GitHub Pages at https://tool-genoa.github.io/genoa-web. The version used in this study is archived on Zenodo at https://doi.org/10.5281/zenodo.18807902 (Wang, 2026). Additional reduction configurations, results, and analyses are provided in the Supplement. Further details are available upon reasonable request.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

Appendix A: Mechanism data structure

A1 Species structure

In GENOA v3, a species (S_i, where $i = 1, \ldots, N_s$ and N_s is the total number of species) is represented as:

$$S_i = \left\{ \begin{array}{l} \text{Name}_i, \text{MW}_i, \text{Formula}_i, \text{SMILES}_i, \text{is\_rdc}, \text{is\_r}, \\ \text{is\_cd}, P_{\text{sat},i}, H_{\text{vap},i}, \text{FG}_{\text{gck, i}}, \text{G}_{\text{gck},i}, \cdots \end{array} \right\} \tag{A1}$$

where Name_i is the species name and also its unique identifier. The core properties of a species used during reduction are listed in Table A1.

Based on stability and volatility, reducible species are classified into four groups:

- Radicals: $S^r = \{S_i \mid \text{is\_r} = 1\}$
- VOCs: $S^{\text{voc}} = \{S_i \mid \text{is\_cd} = 0 \vee P_{\text{sat}} > P_{\text{sat}}^{\text{svoc}}\}$
- SVOCs: $S^{\text{svoc}} = \{S_i \mid \text{is\_cd} = 1 \wedge P_{\text{sat}}^{\text{nvoc}} < P_{\text{sat}} \leq P_{\text{sat}}^{\text{svoc}}\}$
- NVOCs: $S^{\text{nvoc}} = \{S_i \mid \text{is\_cd} = 1 \wedge P_{\text{sat}} \leq P_{\text{sat}}^{\text{nvoc}}\}$

All non-radical groups are restricted to species with $\text{is\_r} = 0$, and species classified as VOCs are updated by setting $\text{is\_cd} = 0$. Here, $P_{\text{sat}}^{\text{svoc}}$ and $P_{\text{sat}}^{\text{nvoc}}$ represent the user-defined thresholds for P_{sat} of semi-volatile and non-volatile organic compounds, respectively. These thresholds can be updated during TBR (Sect. 3.2).

A2 Reaction and kinetic structure

A reaction (R_j, with $j = 1, \ldots, N_r$ and N_r the total number of reactions) is represented as:

$$R_j = \sum_{i \in S_j^{\text{rct}}} \nu_{ij} S_i \rightarrow \sum_{i' \in S_j^{\text{pdt}}} \nu_{i'j} S_{i'}, \tag{A2}$$

where S_j^{rct} and S_j^{pdt} denote the sets of reactants and products for reaction j, and ν_{ij} and $\nu_{i'j}$ are their stoichiometric coefficients. A summary of the core properties associated with the reaction object is provided in Table A2.

Reducible reactions are defined as reactions containing one reducible reactant S_r. For bimolecular reactions, the co-reactant is non-reducible. Under the pseudo-unimolecular assumption, each reducible reaction is expressed as first-order with respect to S_r, while the concentrations of co-reacting reactants are fixed at their reference concentrations. The reaction rate r_j under the selected scenarios is then expressed as:

$$r_j = k_j^{\text{puni}} C_r^{\text{ref}}, \tag{A3}$$

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

Table A1. Core properties of a species object. [a]

Notation	Description	Value [b]
Name	Species name	str
MW	Molecular weight [g mol^{-1}]	flt
Formula	Molecular formula	str
Ratio	Atomic ratios	dict
SMILES	SMILES representation	str
is_rdc	Reducibility indicator	bool
is_r	Radical flag	bool
RO_2	RO_2 group number (0 = non-RO_2)	int
is_cd	Condensability flag	bool
P_{sat}	Saturation vapor pressure [atm]	flt
H_{vap}	Enthalpy of vaporization [kJ mol^{-1}]	flt
FG_{smiles}	Functional structure from SMILES	dict
FG_{gck}	GECKO-A functional group label	str
G_{gck}	GECKO-A generation number	int
d	Node depth	int
S^{up}	Upstream reducible species	set
S^{dwn}	Downstream reducible species	set
S^{prt}	Direct precursor (parent) species	set
S^{cld}	Direct product (child) species	set
C^{ref}	Reference concentration [molec cm^{-3}]	flt
τ	Lifetime [s]	flt
Y^{abs}	absolute effective product yield	flt

[a] Subscript i is omitted for clarity, as all properties refer to a single object.
[b] Value types: str = string, bool = boolean (1/0 for True/False), flt = float, int = integer, dict = key-value map, set = collection of unique elements.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

where the pseudo-unimolecular rate constant k_j^{puni} is defined as:

$$k_j^{\text{puni}} = k_j \prod_{i \in S_j^{\text{rct}} \setminus \{S_r\}} C_i^{\text{ref}}. \tag{A4}$$

Here C_i^{ref} is the reference concentration of the co-reacting non-reducible species S_i.

For example, for a unimolecular reaction $S_r \rightarrow$ products, $k_j^{\text{puni}} = k_j$. For a bimolecular reaction such as $S_r + \text{OH} \rightarrow$ products, the reaction is written as first-order in S_r with $k_j^{\text{puni}} = k_j C_{\text{OH}}^{\text{ref}}$. For RO_2-RO_2 reactions, following the representation used in MCM and GECKO-A, one RO_2 species is treated as reducible, while the other is represented by a non-reducible counter RO_2 class. Its reference concentration is defined as the summed concentration of all RO_2 species in that counter class. Reference concentrations are estimated steady-state values under a predefined set of scenarios and are stored as attributes in the species object (Table A1). These reference concentrations are also used during SBR to calculate lumping ratios.

Table A2. Core properties of reaction and kinetic objects. [a]

Notation	Description	Value[a]
S^{rct}	Set of reactant species	set
S^{pdt}	Set of product species	set
ν	Stoichiometric coefficients	set
$\mathcal{T}$	Reaction type	cat
K	Associated kinetic object	obj
k^{puni}	Pseudo-unimolecular rate constant	flt
r	Pseudo-unimolecular reaction rate	flt
is_rdc	Reducibility indicator	bool
S_r	Targeted reactant species	obj
B	Effective branching ratio	flt

[a] As in Table A1, with additional value types: cat = categorical label, obj = structured object.

A3 Reaction graph structure

In the reaction graph, each reducible species node (S_i with $\text{is\_rdc}_i = 1$) is assigned a node depth d_i, defined as the shortest number of reaction steps from any PVOC species ($S_p \in S^{\text{pvoc}}$).

$$d_i = \min_{p \in S^{\text{pvoc}}} \text{dist}(S_i, S_p). \tag{A5}$$

In addition to node depth, reducible species S_i is associated with graph-based connectivity sets derived from the reaction graph, including:

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

- Downstream species set (S_i^{dwn}): all reducible species reachable from S_i.
- Children species set (S_i^{cld}): species formed directly from the degradation of S_i.
- Upstream species set (S_i^{up}): all species that precede S_i in the graph.
- Parent species set (S_i^{prt}): species that directly produce S_i.

These graph-based sets are also stored in the species object (Table A1).

A species branch rooted at reducible species S_i is defined as the connected subgraph consisting of S_i together with all reducible species and associated reactions downstream of S_i. The branch is denoted S_i^{brh}, with the corresponding species set S_i^{brh}. Branch-level properties are computed by aggregating attributes over S_i^{brh}. Terminal products are defined as reducible species within a branch that have no downstream reducible species in the reaction graph. An example illustrating node depth, connectivity sets, and branch structure is provided in Appendix B5.1.

Appendix B: Additional reduction identification details

B1 TBR strategies and effective yield

Table B1 summarizes the mathematical formulations for TBR strategies.

For the relative and absolute effective product yield strategies, the relative effective product yield of species S_i is computed. Specifically, for a product species S_i generated from reactions that share the same reactants and reaction type, the yield is defined as:

$$\forall\, S_i \in S_p^{\text{cld}}, Y_i^{\text{rel}} = \frac{\sum\limits_{j \in R^*(S_p \to S_i)} k_j^{\text{puni}} \nu_{ij}}{\sum\limits_{S_{i'} \in S_p^{\text{cld}}} \sum\limits_{j' \in R^*(S_p \to S_{i'})} k_{j'}^{\text{puni}} \nu_{i'j'}} \tag{B1}$$

This type-specific relative effective product yield acts as a branching ratio, partitioning product formation among the children of S_p for the same reaction type, with contributions normalized to unity ($\sum_{S_i \in S_p^{\text{cld}}} Y_i^{\text{rel}} = 1$).

Similarly, the overall relative yield $Y_i^{\text{rel,tot}}$ for species S_i across all reaction types can be estimated by replacing the rate constant k^{puni} with the reaction rate r in Eq. B1. With $Y_i^{\text{rel,tot}}$ defined, the absolute effective product yield Y_i^{abs} is then accumulated recursively from the precursor, with $Y_p^{\text{abs}} = 1$ for the initial precursor.

B2 SBR stage, cycle, and step

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.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

Table B1. Mathematical representation of TBR strategies described in Sect. 3.2.

Strategy	Threshold	Mathematical expression [a]
Environmentally insensitive strategy		
NVOC reclassification	$P_{\text{sat}}^{\text{nvoc}}$	$S^{\text{nvoc}} = \{S_i \mid \text{is\_cd}_i = 1 \wedge P_{\text{sat},i} \leq P_{\text{sat}}^{\text{nvoc}}\}$
VOC reclassification	$P_{\text{sat}}^{\text{voc}}$	$S^{\text{voc}} = \{S_i \mid \text{is\_cd}_i = 0 \vee P_{\text{sat},i} > P_{\text{sat}}^{\text{svoc}}\}$
NVOC downstream removal	$P_{\text{sat}}^{\text{nvoc,dwn}}$	$S \setminus \{S_i \in S_b^{\text{dwn}} \mid \exists\, S_b \text{ with } P_{\text{sat},b} < P_{\text{sat}}^{\text{nvoc,dwn}}\}$
VOC branch removal	$P_{\text{sat}}^{\text{svoc,brh}}$	$S \setminus \{S_i \in S_b^{\text{dwn}} \mid \forall\, S_i \in S_b^{\text{brh}},\ P_{\text{sat},i} > P_{\text{sat}}^{\text{svoc,brh}}\}$
Node depth limit	$d^{\max}$	$S \setminus \{S_i \mid d_i < d^{\max}\}$
Relative effective product yield removal	$\left(\frac{Y^{\text{rel}}}{d}\right)_{\min}$	$S \setminus \left\{S_i \mid \frac{Y_i^{\text{rel}}}{d_i} < \left(\frac{Y^{\text{rel}}}{d}\right)_{\min}\right\}$
Fast-degradation lifetime jumping	τ^{fast}	$S \setminus \left\{S_i \mid \tau_i < \tau^{\text{fast}}\right\}$, where $\tau_i = \left(\sum_{j \in R(S_i \to *)} k_j^{\text{puni}}\right)^{-1}$
Environmentally sensitive strategy		
Absolute effective product yield removal	$Y_{\min}^{\text{abs}}$	$S \setminus \left\{S_i \mid Y_i^{\text{abs}} < Y_{\min}^{\text{abs}}\right\}$, where $Y_i^{\text{abs}} = \sum_{S_p \in S_i^{\text{prt}}} Y_p^{\text{abs}} \cdot Y_{p,i}^{\text{rel, tot}}$
Effective reaction branching ratio removal	$B_{\min}$	$R \setminus \{R_j \mid B_j < B_{\min}\}$, where $B_j = \frac{k_j^{\text{puni}}}{\sum_{S_c \in S_i^{\text{cld}}} \sum_{j' \in R(S_i \to S_c)} k_{j'}^{\text{puni}}}$
Branch concentration removal	$C_{\min}^{\text{ref,brh}}$	$S \setminus \left\{S_i^{\text{brh}} \mid \sum_{i' \in S_i^{\text{brh}}} C_{i'}^{\text{ref}} < C_{\min}^{\text{ref,brh}}\right\}$

[a] S and R denote the sets of reducible species and reactions. $S \setminus \ldots$ indicates the removal of species, except in the case of fast-degradation lifetime jumping, where species are bypassed via the jumping strategy. $R \setminus \ldots$ indicates reaction removal.

- 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 becomes empty or no further acceptable reduction is found. The process then continues to the next species group.

- Cycle: A complete pass through all species groups, performing one or more steps per group as needed. At the end of each cycle, an optional evaluation process may be used to assess the reduced mechanism, with results guiding early termination or prompting adjustments to configuration settings.

- Stage: A stage corresponds to a specific configuration defined by the user, including the choice of 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 advances to the next stage. SBR ends once all stages have been completed.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

B3 SBR species grouping and ordering

Species are grouped into non-overlapping sets G_g (with $g = 1, \ldots, N_g$, where N_g is the total number of groups) during the SBR preparation, such that all species within a group share the same values for a user-selected set of attributes:

$$\bigcup_{i=1}^{N_s} S_i = \bigcup_{g=1}^{N_g} G_g \tag{B2}$$

By default, GENOA v3 defines a species group as:

$$G_g = \{S_i \mid S_i \text{ shares Formula, is\_rdc, is\_cd, is\_r, and } \mathrm{G_{gck}}\} \tag{B3}$$

An illustrative example of species grouping is provided in Appendix B5.2.

Each group G_g is assigned a rank R_g that quantifies its relative importance. The rank is computed as the product of three contributions derived from group-level properties:

$$R_g = f_{\mathrm{MW}}(g) \cdot f_d(g) \cdot f_{s \wedge c}(g), \tag{B4}$$

including:

- Mass contribution:

$$f_{\mathrm{MW}}(g) = \frac{|\mathrm{MW}_g - \mathrm{MW}_{S^{\mathrm{pvoc}}}|}{\mathrm{MW}_g + \mathrm{MW}_{S^{\mathrm{pvoc}}}} \tag{B5}$$

  where MW_g is the arithmetic mean molar mass of all species in group, and $\mathrm{MW}_{S^{\mathrm{pvoc}}}$ is that of the PVOC reference set.

- Node depth contribution:

$$f_d(g) = d_g \tag{B6}$$

  where d_g is the mean graph depth of all species in group.

- Stability and condensability contribution:

$$f_{s \wedge c}(g) = \begin{cases} 10 & \text{if is\_r} = 1 \wedge \text{is\_cd} = 0 \\ 1 & \text{if is\_r} = 0 \wedge \text{is\_cd} = 0 \\ 0.1 & \text{if is\_r} = 0 \wedge \text{is\_cd} = 1 \end{cases} \quad \forall S_i \in G_g \tag{B7}$$

  The formulation assumes uniform stability and condensability within each group. Radical species are assigned the highest priority for reduction, followed by stable non-condensable species, and then by condensable stable species.

Groups with higher ranks are considered less critical (or farther from the reference PVOCs) and are therefore processed earlier. Both the grouping and ranking methods can be adjusted to emphasize different reduction priorities.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

Table B2. Mathematical representation of SBR strategies.

Reaction removal	$R \setminus \{R_j \mid R_j \in R(S_p \rightarrow S_i)\}$
Species removal	$S \setminus \{S_i\}$
Gas-particle removal	$S \setminus \{S_i \mid \text{is\_cd}_i > 0 \mapsto \text{is\_cd}_i = 0\}$
Jumping	$\left\{ \begin{array}{l} S_i \mid \ R(S_p \rightarrow S_i) \mapsto R(S_p \rightarrow S_c), \\ S_i, R(S_i \rightarrow S_c) \text{ removed} \end{array} \right\}$
Lumping	$S_i^{\text{lp}} \mapsto S_{\text{new},i}, \ \ f_{\text{lp},i'} = \frac{C_{i'}^{\text{ref}}}{\sum_{S^{\text{lp}}}^{i} C_i^{\text{ref}}}$
Replacement	$S_i^{\text{rp}} \mapsto S_i, \ S_i \in S^{lp}$

B4 SBR candidate identification

Table B2 summarizes the mathematical expressions of the reduction strategies used to identify base candidates for SBR.

A species S_i is eligible for jumping if one or more children $S_c \in S_i^{\text{cld}}$ are identified. The associated reaction $R_j \in R(S_p \rightarrow S_i)$ is then modified as:

$$R_j : S_p \rightarrow \sum_{S_{i'} \in (S_j^{\text{prt}} \setminus S_i)} \nu_{i'j} S_{i'} + \sum_{S_c \in S_i^{\text{cld}}} \left(\nu_{ij} Y_{i,c}^{\text{rel}}\right) S_c, \tag{B8}$$

where species S_i is bypassed from the mechanism, and its role is redistributed to its children via adjusted stoichiometric coefficients. This involves deleting the reaction set $R(S_i \rightarrow S_c)$ and updating $R(S_p \rightarrow S_i)$ to replace S_i with its children S_c.

With the lumping strategy, GENOA defines a lumpable species set S_i^{lp}, which includes S_i from the targeted species group and other species with similar properties (Table B3) that are not connected in the reaction graph. A new surrogate species $S_{\text{new,i}}$ is then constructed from this set. For each $S_{i'} \in S_i^{lp}$, a lumping ratio is calculated as:

$$f_{\text{lp},i'} = \frac{C_{i'}^{\text{ref}}}{\sum_{S_k \in S_i^{lp}}^{k} C_k^{\text{ref}}}. \tag{B9}$$

This ratio quantifies the relative contribution of each species within the group. In practice, all species in S_i^{lp} are replaced by a surrogate species $S_{\text{new},i}$. The surrogate structure is constructed from the species with the highest $f_{\text{lp},i'}$ and is used to define the representative molecular structure in the reaction graph. The physical and chemical properties of $S_{\text{new},i}$ are then updated using weighted aggregation of all species in S_i^{lp} based on the lumping ratios (Table B4).

The replacement strategy can be regarded as a simplified form of lumping, as both search for similar species, but replacement does not construct a new surrogate. Instead, species that are co-products of the same reactions and meet the similarity criteria in Table B3 are grouped into a replaceable set. Depending on the settings, any species in the set may be selected as the replacer, with either one species substituting all others or only a subset being replaced.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

Table B3. Criteria for identifying similar species during reductions via lumping or replacement.

Prop. [a]	Criterion for similarity
Lumping	
is_r	Same stability flag
is_cd	Same condensability flag
MW	Absolute diff. $\lvert \mathrm{MW}_i - \mathrm{MW}_{i'} \rvert \leq 20$ g mol^{-1}
MW	Relative diff. $\frac{\lvert \mathrm{MW}_i - \mathrm{MW}_{i'} \rvert}{(\mathrm{MW}_i + \mathrm{MW}_{i'})/2} \leq 0.1$
$\mathrm{FG}_{\mathrm{gck}}$	Same GECKO-A functional group
P_{sat}	$\lvert \log_{10}(P_{\mathrm{sat},i}) - \log_{10}(P_{\mathrm{sat},i'}) \rvert \leq 2$
τ	$0.1 \leq \tau_i / \tau_{i'} \leq 10$
nC	Diff. in carbon count ≤ 3
nO	Diff. in oxygen group count ≤ 3
Replacement	
MW	Relative diff. $\frac{\lvert \mathrm{MW}_i - \mathrm{MW}_{i'} \rvert}{(\mathrm{MW}_i + \mathrm{MW}_{i'})/2} \leq 0.5$
nC	Diff. in carbon count ≤ 3
link	Co-products of the same reactions

[a] Property notation follows Table A1. S_i denotes the target species, and $S_{i'}$ the candidate being tested for inclusion in the lumpable (S_i^{lp}) or replaceable (S_i^{rp}) sets. nC and nO are the numbers of carbon and oxygen atoms derived from $\mathrm{FG}_{\mathrm{smiles}}$; Diff. = Difference; link = graph-based connectivity/relationship.

B5 Reduction demonstration using a simple mechanism

This example walks through mechanism preparation and the SBR steps using an illustrative mechanism. The chemistry shown here is intentionally simplified and not intended to be atmospherically realistic.

B5.1 Example mechanism and preparation

We consider a small degradation scheme consisting of seven reactions from a PVOC, denoted as $\mathrm{R_{org}}1$ to $\mathrm{R_{org}}7$ in Table B5. Further degradation beyond these seven reactions is ignored for simplicity. For the targeted portion of the mechanism, the reducible species are:

$$\{\text{RO2-1, RO-1, VOC-1, SVOC-1, RO2-2, SVOC-2, VOC-2}\}. \tag{B10}$$

Following the preparation described in Sect. 4.1, seven original reactions are merged into five pre-processed reactions (R1–R5; Table B5), with $\mathrm{R_{org}}1$–$\mathrm{R_{org}}2$ merged as R1 and $\mathrm{R_{org}}3$–$\mathrm{R_{org}}4$ as R2. The resulting mechanism is shown as a reaction graph in Fig. B1.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

Table B4. Lumped surrogate properties for S_{new}.

Prop.	Method [a]	Expression
Name	Dom.	From species with highest f_{lp}
is_cd	Dom.	Shared by all lumped species
is_r	Dom.	Shared by all lumped species
RO_2	Dom.	From species with highest f_{lp}
FG_{gck}	Dom.	Shared by all lumped species
G_{gck}	Dom.	From species with highest f_{lp}
MW	Avg.	$\sum f_{\text{lp}} \cdot MW$
FG_{smile}	Avg.	Elemental mean
ratio	Recomp.	$\sum f_{\text{lp}} \cdot \text{Ratio}$
C^{ref}	Sum	$\sum C^{\text{ref}}$
P_{sat}	Geo.	$10^{\sum f_{\text{lp}} \cdot \log_{10}(P_{\text{sat}})}$
H_{vap}	Avg.	$\sum f_{\text{lp}} \cdot H_{\text{vap}}$
d	Recomp.	From updated reaction graph
S^{up}	Recomp.	From updated reaction graph
S^{dwn}	Recomp.	From updated reaction graph
S^{prt}	Recomp.	From updated reaction graph
S^{cld}	Recomp.	From updated reaction graph
Y^{abs}	Recomp.	From updated reaction graph

[a] Method abbreviations: Dom. = dominant species with the highest lumping ratio f_{lp}; Avg. = weighted arithmetic mean; Geo. = weighted geometric mean (log space); Recomp. = recomputed based on updated graph or values; Sum = direct summation.

In this reaction graph, PVOC is taken as the branch root with node depth $d = 0$ (Eq. A5). RO2-1 and RO2-2 have $d = 1$, VOC-1 has $d = 3$, and the remaining reducible species have $d = 2$, giving a maximum node depth of 3. Four species (VOC-1, SVOC-1, VOC-2, and SVOC-2) are considered terminal products, as no further reactions leading to other reducible species originate from them. The graph-based connectivity sets (Appendix A3) can also be illustrated. For example,

- RO2-1 has RO-1 and SVOC-1 as its children species set; together with VOC-1, they form its downstream species set.
- For VOC-1, its parent set contains RO-1, while its upstream species set includes PVOC and RO2-1 in addition to RO-1.

Tracer species are also included in the pre-processed mechanism to quantify the loss or production of selected compounds (OH, NO, NO_2, HO_2). These tracers are used as evaluation targets for SBR (Sect. 4.4). In this example, lOH and lNO track the loss of OH and NO attributable to the VOC degradation scheme, while pNO_2 and pHO_2 track the corresponding production. After each reduction attempt, tracer species are updated to preserve the net stoichiometric budgets of the tracked species.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

This is illustrated by reaction R′1 in Table B5, which is obtained by jumping RO2-1 and combining R1 and R2, resulting in redistributed but conserved tracer contributions.

Table B5. Example reaction lists showing original reactions, pre-processed reactions, and the combined reaction R′1 after jumping RO2-1. The corresponding reaction graph for the pre-processed reactions is shown in Fig. B1.

Original reactions (R_{org}, N = 7)	
$R_{org}1$	PVOC + OH → RO2-1
$R_{org}2$	PVOC + OH → RO2-2
$R_{org}3$	RO2-1 + NO → RO-1 + NO_2
$R_{org}4$	RO2-1 + NO → SVOC-1
$R_{org}5$	RO-1 → VOC-1 + HO_2
$R_{org}6$	RO2-2 + HO_2 → SVOC-2
$R_{org}7$ [a]	RO2-2 + RO_2 → VOC-2
Pre-processed reactions with tracers [b] (R, N = 5)	
R1 [c]	PVOC + OH → $f_{R1,1}$·RO2-1 + $f_{R1,2}$·RO2-2 + lOH ($R_{org}1$ + $R_{org}2$ ⇒ R1 with k1)
R2	RO2-1 + NO → $f_{R2,1}$·RO-1 + $f_{R2,1}$·NO_2 + $f_{R2,2}$·SVOC-1 + lNO + pNO_2 ($R_{org}3$ + $R_{org}4$ ⇒ R2 with k2)
R3	RO-1 → VOC-1 + HO_2 + pHO_2 ($R_{org}5$ ⇒ R3 with k3)
R4	RO2-2 + HO_2 → SVOC-2 + lHO_2 ($R_{org}6$ ⇒ R4 with k4)
R5	RO2-2 + RO_2 → VOC-2 ($R_{org}7$ ⇒ R5 with k5)
Reactions R1 and R2 after jumping RO2-1	
R'1	PVOC + OH → $f_{R1,1}$ · $f_{R2,1}$· RO-1 + $f_{R1,1}$ · $f_{R2,1}$· NO_2 + $f_{R1,1}$ · $f_{R2,2}$· SVOC-1 + $f_{R1,1}$ · $f_{R2,1}$· pNO_2 + $f_{R1,2}$·RO2-2 + lOH (R1 + R2 ⇒ R′1 with k1)

[a] RO_2 + RO_2 permutation reaction with RO_2 representing the total RO_2 class (RO2-1 + RO2-2).
[b] Loss and production of OH, NO, NO_2, and HO_2 are tracked using tracer species prefixed with l and p, respectively.
[c] Stoichiometric coefficients $f_{Ri,j}$ are placeholders and are not explicitly computed in this example.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

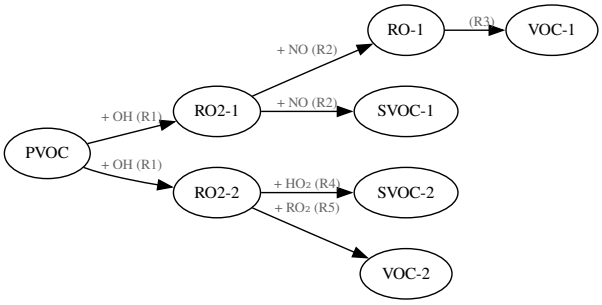

Figure B1. The reaction graph representing the pre-processed example mechanism (R1 to R5). PVOC and reducible species are shown in circled nodes, reactions are shown as arrows pointing from reactants to products with non-reducible reactants and reaction identifiers noted above.

B5.2 SBR species grouping

After preparation, reducible species are grouped prior to SBR. The grouping methodology and ranking equations are described in Appendix B3. Species are grouped based on chemical stability class (for example, peroxy radicals, alkoxy radicals, or stable species) and volatility. A hypothetical grouping for this example is:

$$\begin{aligned} G_1 &= \{\text{RO2-1}, \text{RO2-2}\}, \\ G_2 &= \{\text{RO-1}\}, \\ G_3 &= \{\text{SVOC-1}, \text{SVOC-2}\}, \\ G_4 &= \{\text{VOC-1}, \text{VOC-2}\}. \end{aligned} \tag{B11}$$

The groups are then processed sequentially according to the ranking metric. Each group is processed independently to generate and evaluate reduction candidates.

B5.3 SBR candidate identification

For each species group, GENOA v3 identifies base and variant reduction candidates using the SBR strategies (Sect. B4). Candidate identification is illustrated for group G_1 at one SBR step (Appendix B2).

- Species removal: The base candidate removes all species in the group, while variants remove subsets of the group. For G_1, the base candidate removes RO2-1 and RO2-2, and variants remove either RO2-1 or RO2-2 only.
- Reaction removal: Reactions associated with a group include those involving group species as reactants or terminal products (i.e., with no subsequent reactions). The base candidate removes all such reactions, while variants remove subsets. For G_1, the associated reactions are R2, R4, and R5, which are removed individually or in combination (one base candidate and six variants).

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

- Gas–particle partitioning removal: This strategy disables gas–particle partitioning while retaining gas-phase chemistry and applies only to condensable species. It does not apply to G_1, which contains only radicals, but is applied to G_3, where partitioning is disabled for one or both species.

- Lumping: Similarity checks (Table B3) are applied to identify lumpable species, either within the group or across all reducible species. For G_1, RO2-1 and RO2-2 are assumed to satisfy lumping criteria and are lumped into a new surrogate, with properties from Table B3 and lumping ratios computed using Eq. B9. No variant candidates exist for G_1, but variants can be generated using alternative lumping ratios or, when multiple lumpable sets exist or a set contains more than two species, by evaluating subsets with redistributed ratios.

- Replacement: Replacement applies to co-products of the same reaction that satisfy the similarity checks (Table B3). For G_1, RO2-1 and RO2-2 are co-products of reaction R1 and are assumed to be similar, resulting in two base candidates (RO2-1 replaced by RO2-2, or vice versa). No variant candidates exist for G_1, but variants can be generated when multiple replacement sets exist or when a set contains more than two species.

- Jumping: Jumping applies to species with a single destruction pathway. For G_1, RO2-1 is eligible and reactions R1 and R2 are combined into R′1 after jumping (Table B5), while RO2-2 is not eligible due to multiple degradation pathways (R4 and R5). No variant candidates exist for G_1, but variants can be generated when multiple jumping sets are identified.

In total, six base candidates and eight variant candidates are generated for G_1 and following the dynamic procedure described in Sect. 3.3.3 and Appendix C5, with base candidates examined before variants. Once a candidate is accepted, the affected species groups are updated and the step is complete.

For example, removal of reaction R5 updates G_4 to $G_4^{\text{new}} = \{\text{VOC-1}\}$ (no changed in G_1), while RO2-1 jumping updates G_1 to $G_1^{\text{new}} = \{\text{RO2-2}\}$. In subsequent steps, candidate identification continues for G_1 (or G_1^{new}) until it is empty or no further reduction is possible, then proceeds to the next species group.

After all species groups have been processed (completion of one SBR cycle), the full species group list and their ordering are recomputed based on the updated mechanism.

Appendix C: Additional reduction evaluation details

C1 Size reduction quantification

Mechanism size reduction can be quantified using the retained size factor (f_{ret}):

$$f_{\text{ret}} = \frac{N^{\text{rn}}_{\text{rdc}} + N^{\text{sp}}_{\text{rdc}} + N^{\text{cd}}_{\text{rdc}}}{N^{\text{rn}}_{\text{ref}} + N^{\text{sp}}_{\text{ref}} + N^{\text{cd}}_{\text{ref}}} \times 100\%, \tag{C1}$$

where N^{rn}, N^{sp}, and N^{cd} represent the numbers of reducible reactions, species, and condensable species, respectively. Subscripts “ref” and “rdc” refer to the reference (unreduced) and reduced mechanisms. Smaller f_{ret} correspond to stronger reductions. The complementary metric, relative size reduction factor (f_{rdc}), is defined as

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

$$f_{\text{rdc}} = 100\ \% - f_{\text{rdc}}, \tag{C2}$$

which represents the percentage decrease relative to the reference mechanism size and is used throughout the results and discussion sections.

During SBR candidate evaluation, reduction efficiency is quantified by a weighted reduction score for each candidate, defined as:

$$O = f_{\text{rn}}\Delta N^{\text{rn}} + f_{\text{sp}}\Delta N^{\text{sp}} + f_{\text{cd}}\Delta N^{\text{cd}}, \tag{C3}$$

where ΔN^{rn}, ΔN^{sp}, and ΔN^{cd} are the changes in the numbers of reactions, species, and condensable species, respectively. The weights f_{rn}, f_{sp}, and f_{cd} control the relative importance of each component in the score and are user-adjustable. The test cases in this study use $f_{\text{rn}} = 1$, $f_{\text{sp}} = 1$, and $f_{\text{cd}} = 10$, consistent with the settings of the previous version.

C2 Reduction error for a simulation

Consider a simulation L_l (with $l = 1, \ldots, N_L$), corresponding to one evaluation scenario with a two-day simulation period, as used in this work. For each evaluation target S_s^{eval} (with $s = 1, \ldots, N_{S^e}$) and reference mechanism M_r^{ref} (with $r = 1, \ldots, N_{M^r}$), the reduction error over a time segment (T_d) is defined using the fractional absolute error (FAE):

$$\epsilon_{s,r,d} = 2\,\frac{\sum_{t\in T_d} \left|C_s(t) - C_{s,r}^{\text{ref}}(t)\right|}{\sum_{t\in T_d} \left|C_s(t) + C_{s,r}^{\text{ref}}(t)\right|}, \tag{C4}$$

where $C_s(t)$ and $C_{s,r}^{\text{ref}}(t)$ are time-dependent concentrations predicted by the reduced and reference mechanisms, respectively. Here, d indexes the time segments, with $d \in \{1, 2\}$ corresponding to day 1 and day 2 in the two-day simulation.

The reduction error for simulation L_l is defined as the maximum over all time segments:

$$E_{s,r}^{l} = \max_{d\in\{1,2\}} \epsilon_{s,r,d}. \tag{C5}$$

All aspects of the error definition, including the number of reference mechanisms and evaluation targets, error metric, time segmentation, and aggregation method (e.g., maximum over segments), can be adjusted depending on the application.

C3 Reduction error for a candidate/mechanism

A reduced mechanism can be evaluated over N_L simulations, each corresponding to an evaluation scenario. When associated with a reduction candidate χ_k (with $k = 1, \ldots, N_{\text{rdc}}$) during SBR, three error metrics are used for SBR candidate evaluation:

- Mean reduction error:

$$\overline{E_k} = \frac{1}{N_L\ N_{S^e}\ N_{M^r}} \sum_{l=1}^{N_L}\sum_{s=1}^{N_{S^e}}\sum_{r=1}^{N_{M^r}} E_{s,r,k}^{l}; \tag{C6}$$

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

- Maximum reduction error:

$$E_k^{\max} = \max_{\substack{l=1,\dots,N_L \\ s=1,\dots,N_{S^e} \\ r=1,\dots,N_{M^r}}} E_{s,r,k}^{l}; \tag{C7}$$

- Reduction error increase:

$$\Delta\overline{E_k} = \max(\overline{E_k} - \overline{E_{\text{pre}}},\ E_k^{\max} - E_{\text{pre}}^{\max}), \tag{C8}$$

where $\overline{E_{\text{pre}}}$ and $E_{\text{pre}}^{\max}$ denote the mean and maximum reduction errors obtained by the last accepted reduction candidate.

When used to assess the accuracy of a reduced mechanism (e.g., Figs. 3 and 6), the reduction error refers to the mean reduction error ($\overline{E}$), averaged over the evaluation targets and scenarios specified in the main text (by default, across 11 scenarios).

C4 Reduction evaluation metrics

Predefined error tolerance sets are applied to SBR candidate evaluation. Three error tolerances are currently applied, corresponding to the mean error, maximum error, and the increase in both mean and maximum errors, respectively. A reduction candidate χ_k is acceptable if:

$$\overline{E_k} \le \epsilon_{\text{avg}} \ \wedge\ E_k^{\max} \le \epsilon_{\max} \ \wedge\ \Delta\overline{E_k} \le \epsilon_{\Delta}, \tag{C9}$$

where ϵ_{avg}, $\epsilon_{\max}$, and ϵ_{Δ} denote the error tolerances for mean error, maximum error, and error increases relative to the previous mechanism for both mean and maximum errors, respectively.

Acceptable candidates are those that meet all tolerances (χ^{acc}). The optimal reduction candidate χ^{opt} is then determined as the acceptable candidate with the highest reduction score (Eq. C3):

$$\chi^{\text{opt}} = \arg\max_{k \in \chi^{\text{acc}}} O_k. \tag{C10}$$

If a tie occurs, the candidate with the lowest $\overline{E_k}$ among those with the same score is selected. In the absence of further restrictions, the optimal candidate is accepted as the reduced mechanism and serves as the basis for subsequent reduction steps, with the errors of the last accepted candidate updated as:

$$\overline{E_{\text{pre}}} = \overline{E}_{\chi^{\text{opt}}};\quad E_{\text{pre}}^{\max} = E_{\chi^{\text{opt}}}^{\max}. \tag{C11}$$

C5 Dynamic evaluation

A full evaluation of all candidates would require $N_l \times N_{\text{rdc}}$ simulations, which is computationally demanding, particularly in the early stages when mechanisms remain large. To accelerate the process, GENOA v3 applies dynamic evaluation, in which candidates are ranked by their reduction scores before simulations are launched.

Candidates (χ_k) with the same highest score (O') are grouped:

$$\mathcal{G}(O') = \{\chi_k \mid O_k = O'\} \tag{C12}$$

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

Simulations for $\mathcal{G}(O')$ are launched in parallel. As soon as a simulation finishes, error checks are performed. Candidates exceeding the tolerances for maximum error or error increase are immediately rejected, and their remaining simulations are canceled. For the surviving candidates, mean error tolerances are evaluated once all simulations are completed.

If one or more candidates in $\mathcal{G}(O')$ satisfy all tolerances, the candidate χ^* with the smallest $\overline{E_k}$ is selected and accepted for the reduction, while lower-scoring candidates are skipped. If no candidate in $\mathcal{G}(O')$ passes, the next score group is considered.

This dynamic evaluation process identifies the optimal candidate while avoiding full simulation of all candidates, thereby minimizing computational cost and improving efficiency.

Author contributions. ZW led conceptualization, methodology development, software development and implementation, performance analysis and model validation, visualization, resource and manuscript preparation; WPLC contributed to methodology and analysis, and writing review and editing; KB led funding acquisition, project administration, and supervision, contributed to methodology, writing the initial draft, and review and editing; RV, MC, BA, and JLT provided GECKO-A resources and coupling support, contributed to analysis and writing review and editing; JO and QY contributed to analysis and interpretation, and writing review and editing.

Competing interests. The authors declare that no competing interests are present.

Acknowledgements. This research was funded by the U.S. Environmental Protection Agency (EPA) through the Science to Achieve Results (STAR) program under Agreement No. 84000701.

Disclaimer. This publication has not been formally reviewed by the U.S. Environmental Protection Agency (EPA). The views expressed are those of the authors and do not necessarily represent the views or policies of the EPA. The EPA does not endorse any products or commercial services mentioned in this publication.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

References

Atkinson, R. and Arey, J.: Atmospheric Degradation of Volatile Organic Compounds, Chem. Rev., 103, 4605–4638, https://doi.org/10.1021/cr0206420, pMID: 14664626, 2003.

Aumont, B., Szopa, S., and Madronich, S.: Modelling the evolution of organic carbon during its gas-phase tropospheric oxidation: development of an explicit model based on a self generating approach, Atmos. Chem. Phys., 5, 2497–2517, https://doi.org/10.5194/acp-5-2497-2005, 2005.

Aumont, B., Camredon, M., Lee-Taylor, J., and Valorso, R.: Box Model for GECKO-A, https://doi.org/10.5281/zenodo.15310153, 2025a.

Aumont, B., Camredon, M., Lee-Taylor, J., and Valorso, R.: GECKO-A: Generator for Explicit Chemistry and Kinetics of Organics in the Atmosphere, https://doi.org/10.5281/zenodo.15309905, 2025b.

Bianchi, F., Kurtén, T., Riva, M., Mohr, C., Rissanen, M. P., Roldin, P., Berndt, T., Crounse, J. D., Wennberg, P. O., Mentel, T. F., Wildt, J., Junninen, H., Jokinen, T., Kulmala, M., Worsnop, D. R., Thornton, J. A., Donahue, N., Kjaergaard, H. G., and Ehn, M.: Highly Oxygenated Organic Molecules (HOM) from Gas-Phase Autoxidation Involving Peroxy Radicals: A Key Contributor to Atmospheric Aerosol, Chem. Rev., 119, 3472–3509, https://doi.org/10.1021/acs.chemrev.8b00395, 2019.

Camredon, M., Aumont, B., Lee-Taylor, J., and Madronich, S.: The SOA/VOC/NO x system: an explicit model of secondary organic aerosol formation, Atmos. Chem. Phys., 7, 5599–5610, https://doi.org/10.5194/acp-7-5599-2007, 2007.

Carter, W. P. L.: Documentation of the SAPRC-99 Chemical Mechanism for VOC Reactivity Assessment, https://doi.org/10.5281/zenodo.12600705, 2000.

Carter, W. P. L.: Estimation of Rate Constants for Reactions of Organic Compounds under Atmospheric Conditions, Atmosphere, 12, https://doi.org/10.3390/atmos12101250, 2021.

Carter, W. P. L., Jiang, J., Orlando, J. J., and Barsanti, K. C.: Derivation of atmospheric reaction mechanisms for volatile organic compounds by the SAPRC mechanism generation system (MechGen), Atmos. Chem. Phys., 25, 199–242, https://doi.org/10.5194/acp-25-199-2025, 2025a.

Carter, W. P. L., Jiang, J., Wang, Z., and Barsanti, K. C.: The SAPRC Atmospheric Chemical Mechanism Generation System (MechGen), EGUsphere, 2025, 1–37, https://doi.org/10.5194/egusphere-2025-1183, 2025b.

Ehn, M., Thornton, J. A., Kleist, E., Sipilä, M., Junninen, H., Pullinen, I., Springer, M., Rubach, F., Tillmann, R., Lee, B., et al.: A large source of low-volatility secondary organic aerosol, Nature, 506, 476–479, https://doi.org/10.1038/nature13032, 2014.

Ervens, B., Rickard, A., Aumont, B., Carter, W. P. L., McGillen, M., Mellouki, A., Orlando, J., Picquet-Varrault, B., Seakins, P., Stockwell, W. R., Vereecken, L., and Wallington, T. J.: Opinion: Challenges and needs of tropospheric chemical mechanism development, Atmos. Chem. Phys., 24, 13 317–13 339, https://doi.org/10.5194/acp-24-13317-2024, 2024.

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, https://doi.org/10.1051/0004-6361/202039956, 2022.

Hallquist, M., Wenger, J. C., Baltensperger, U., Rudich, Y., Simpson, D., Claeys, M., Dommen, J., Donahue, N. M., George, C., Goldstein, A. H., Hamilton, J. F., Herrmann, H., Hoffmann, T., Iinuma, Y., Jang, M., Jenkin, M. E., Jimenez, J. L., Kiendler-Scharr, A., Maenhaut, W., McFiggans, G., Mentel, T. F., Monod, A., Prévôt, A. S. H., Seinfeld, J. H., Surratt, J. D., Szmigielski, R., and Wildt, J.: The formation, properties and impact of secondary organic aerosol: current and emerging issues, Atmos. Chem. Phys., 9, 5155–5236, https://doi.org/10.5194/acp-9-5155-2009, 2009.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

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, https://doi.org/10.1016/j.atmosenv.2008.07.028, 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, https://doi.org/10.1016/S1352-2310(96)00105-7, 1997.

Jenkin, M. E., Valorso, R., Aumont, B., Rickard, A. R., and Wallington, T. J.: Estimation of rate coefficients and branching ratios for gas-phase reactions of OH with aliphatic organic compounds for use in automated mechanism construction, Atmos. Chem. Phys., 18, 9297–9328, https://doi.org/10.5194/acp-18-9297-2018, 2018.

Jenkin, M. E., Valorso, R., Aumont, B., Newland, M. J., and Rickard, A. R.: Estimation of rate coefficients for the reactions of O_3 with unsaturated organic compounds for use in automated mechanism construction, Atmos. Chem. Phys., 20, 12 921–12 937, https://doi.org/10.5194/acp-20-12921-2020, 2020.

Kaiser, J., Skog, K. M., Baumann, K., Bertman, S. B., Brown, S. B., Brune, W. H., Crounse, J. D., de Gouw, J. A., Edgerton, E. S., Feiner, P. A., Goldstein, A. H., Koss, A., Misztal, P. K., Nguyen, T. B., Olson, K. F., St. Clair, J. M., Teng, A. P., Toma, S., Wennberg, P. O., Wild, R. J., Zhang, L., and Keutsch, F. N.: Speciation of OH reactivity above the canopy of an isoprene-dominated forest, Atmos. Chem. Phys., 16, 9349–9359, https://doi.org/10.5194/acp-16-9349-2016, 2016.

Kanakidou, M., Seinfeld, J., Pandis, S., Barnes, I., Dentener, F. J., Facchini, M. C., Dingenen, R. V., Ervens, B., Nenes, A., Nielsen, C., et al.: Organic aerosol and global climate modelling: a review, Atmos. Chem. Phys., 5, 1053–1123, https://doi.org/10.5194/acp-5-1053-2005, 2005.

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, https://doi.org/10.1029/2021MS002926, 2022.

Kerdouci, J., Picquet-Varrault, B., and Doussin, J.-F.: Structure–activity relationship for the gas-phase reactions of NO3 radical with organic compounds: Update and extension to aldehydes, Atmos. Environ., 84, 363–372, https://doi.org/https://doi.org/10.1016/j.atmosenv.2013.11.024, 2014.

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, https://doi.org/10.5194/acp-18-13411-2018, 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, https://doi.org/10.5194/acp-20-4905-2020, 2020.

Lee-Taylor, J. and Drews, C.: GECKO-A Output Library: Standard Reference Configurations, https://www.acom.ucar.edu/gecko/output-library.shtml, accessed 2026-02-07, 2024.

Lelieveld, J., Evans, J. S., Fnais, M., Giannadaki, D., and Pozzer, A.: The contribution of outdoor air pollution sources to premature mortality on a global scale, Nature, 525, 367–371, https://doi.org/10.1038/nature15371, 2015.

Li, E. W., Sturm, P. O., Silva, S. J., Barber, V. A., and Keller, C. A.: Characterizing the Speed of Chemical Cycling in the Atmosphere, Geophys. Res. Lett., 52, e2024GL111 021, https://doi.org/10.1029/2024GL111021, 2025.

Mouchel-Vallon, C. and Hodzic, A.: Toward emulating an explicit organic chemistry mechanism with random forest models, J. Geophys. Res.-Atmos., 128, e2022JD038 227, https://doi.org/10.1029/2022JD038227, 2023.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

Mouchel-Vallon, C., Lee-Taylor, J., Hodzic, A., Artaxo, P., Aumont, B., Camredon, M., Gurarie, D., Jimenez, J.-L., Lenschow, D. H., Martin, S. T., et al.: Exploration of oxidative chemistry and secondary organic aerosol formation in the Amazon during the wet season: explicit modeling of the Manaus urban plume with GECKO-A, Atmos. Chem. Phys., 20, 5995–6014, https://doi.org/10.5194/acp-20-5995-2020, 2020.

Nannoolal, Y., Rarey, J., Ramjugernath, D., and Cordes, W.: Estimation of pure component properties: Part 1. Estimation of the normal boiling point of non-electrolyte organic compounds via group contributions and group interactions, Fluid Phase Equilibr., 226, 45–63, https://doi.org/10.1016/j.fluid.2004.09.001, 2004.

Nannoolal, Y., Rarey, J., and Ramjugernath, D.: Estimation of pure component properties: Part 3. Estimation of the vapor pressure of non-electrolyte organic compounds via group contributions and group interactions, Fluid Phase Equilibr., 269, 117–133, https://doi.org/10.1016/j.fluid.2008.04.020, 2008.

Newland, M. J., Mouchel-Vallon, C., Valorso, R., Aumont, B., Vereecken, L., Jenkin, M. E., and Rickard, A. R.: Estimation of mechanistic parameters in the gas-phase reactions of ozone with alkenes for use in automated mechanism construction, Atmos. Chem. Phys., 22, 6167–6195, https://doi.org/10.5194/acp-22-6167-2022, 2022.

Orlando, J. J. and Tyndall, G. S.: Laboratory studies of organic peroxy radical chemistry: an overview with emphasis on recent issues of atmospheric significance, Chem. Soc. Rev., 41, 6294–6317, https://doi.org/10.1039/C2CS35166H, 2012.

Orlando, J. J., Tyndall, G. S., and Wallington, T. J.: The Atmospheric Chemistry of Alkoxy Radicals, Chem. Rev., 103, 4657–4690, https://doi.org/10.1021/cr020527p, pMID: 14664628, 2003.

Porter, W. C., Jimenez, J. L., and Barsanti, K. C.: Quantifying atmospheric parameter ranges for ambient secondary organic aerosol formation, ACS Earth Space Chem., 5, 2380–2397, https://doi.org/10.1021/acsearthspacechem.1c00090, 2021.

Roldin, P., Ehn, M., Kurtén, T., Olenius, T., Rissanen, M. P., Sarnela, N., Elm, J., Rantala, P., Hao, L., Hyttinen, N., et al.: The role of highly oxygenated organic molecules in the Boreal aerosol-cloud-climate system, Nat. Commun., 10, 1–15, https://doi.org/10.1038/s41467-019-12338-8, 2019.

Sartelet, K., Couvidat, F., Wang, Z., Flageul, C., and Kim, Y.: SSH-Aerosol v1. 1: A Modular Box Model to Simulate the Evolution of Primary and Secondary Aerosols, Atmosphere, 11, 525, https://doi.org/10.3390/atmos11050525, 2020.

Sartelet, K., Wang, Z., Lannuque, V., Iyer, S., Couvidat, F., and Sarica, T.: Modelling molecular composition of SOA from toluene photooxidation at urban and street scales, Environ. sci.: Atmos., 4, 839–847, https://doi.org/10.1039/D4EA00049H, 2024.

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, https://doi.org/10.5194/acp-3-161-2003, 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, https://doi.org/10.1029/2021MS002974, 2022.

Seinfeld, J. H. and Pandis, S. N.: Atmospheric chemistry and physics: from air pollution to climate change, John Wiley & Sons, 2016.

Silva, S. J. and Halappanavar, M. M.: Graph characterization of higher-order structure in atmospheric chemical reaction mechanisms, Environ. Data Sci., 3, e20, https://doi.org/10.1017/eds.2024.30, 2024.

Stockwell, W. R., Lawson, C. V., Saunders, E., and Goliff, W. S.: A Review of Tropospheric Atmospheric Chemistry and Gas-Phase Chemical Mechanisms for Air Quality Modeling, Atmosphere, 3, 1–32, https://doi.org/10.3390/atmos3010001, 2012.

https://doi.org/10.5194/egusphere-2026-1921
Preprint. Discussion started: 30 April 2026

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, https://doi.org/10.5194/acp-5-2519-2005, 2005.

Topping, D., Barley, M., Bane, M. K., Higham, N., Aumont, B., Dingle, N., and McFiggans, G.: UManSysProp v1.0: an online and open-source facility for molecular property prediction and atmospheric aerosol calculations, Geosci. Model Dev., 9, 899–914, https://doi.org/10.5194/gmd-9-899-2016, 2016.

Vereecken, L. and Francisco, J. S.: Theoretical studies of atmospheric reaction mechanisms in the troposphere, Chem. Soc. Rev., 41, 6259–6293, https://doi.org/10.1039/C2CS35070J, 2012.

Wang, Z.: GENOA v3.0.0, https://doi.org/10.5281/zenodo.18807902, 2026.

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, https://doi.org/10.5194/gmd-15-8957-2022, 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. Atmos. Sci., 174, 106 248, https://doi.org/10.1016/j.jaerosci.2023.106248, 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, https://doi.org/10.1016/j.scitotenv.2024.172142, 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, https://doi.org/10.5194/acp-20-10889-2020, 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, https://doi.org/10.5194/acp-4-2025-2004, 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, https://doi.org/10.5194/acp-4-2057-2004, 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, https://doi.org/10.5194/gmd-16-1801-2023, 2023.

Wiser, F., Sen, S., Wang, Z., Lee-Taylor, J., Barsanti, K., Orlando, J., Westervelt, D., Henze, D., Fiore, A., Berman, A., et al.: A Graph Theory-Based Algorithm for the Reduction of Atmospheric Chemical Mechanisms, ChemRxiv, https://doi.org/10.26434/chemrxiv-2025-5s46n, 2025.

Wolfe, G. M., Marvin, M. R., Roberts, S. J., Travis, K. R., and Liao, J.: The Framework for 0-D Atmospheric Modeling (F0AM) v3.1, Geosci. Model Dev., 9, 3309–3319, https://doi.org/10.5194/gmd-9-3309-2016, 2016.

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, https://doi.org/10.1029/2024JH000132, 2024.