

This is a limited, informal, admittedly opinionated list intended to address modeling pitfalls. Many may seem obvious, but to my intended audience of practical modelers for realistic applications, I have seen these repeatedly over the many years I have been involved, and I continue to see them. Space and clarity preclude me from extensive qualifications of my remarks, which I must trust the reader to infer. I hope you will find something here to help you in your work.

Chemistry First. An aftertreatment reactor (ATR) is, first and foremost, a chemical reactor. No matter how much you know and love transport, heat transfer, and so forth, if you lose focus on the chemistry, you risk missing the main conclusions. Have your organization mix in some chemical engineers with the mechanical engineers who often do the ATR modeling.

Kinetics Rule. Moreover, the chemistry is usually the weak link in our modeling and so provides the lower bound on the accuracy of our predictions. Therefore, why waste time chasing down the effect of secondary processes to the third decimal place, when the kinetic parameters are only known to 5%? Those with experience in obtaining kinetic parameters directly from dedicated experiments will admit that 5% is sometimes optimistic, regardless of the number of digits you see in the literature. Regarding claims that we are on the verge of a golden age when kinetic parameters just fall out from elementary steps and first principles, consider that such claims have been made for decades. If you are tempted to believe such a claim by extrapolating from homogeneous gas-phase reactions, I suggest you read more of the surface chemistry literature to help appreciate the complexities of heterogeneous catalysis.

Modeling is not simulation. If you want to see the effect of every last detail, do an experiment. If you want to see the effect of the main processes, unencumbered by less important or poorly understood details, you model. These are two distinct goals, and you better choose before you start. If you are simulating to replace an experiment, then you have to restrain your expectations to match the limitations of the kinetics (order of 5%, above). Embrace modeling, rather than simulation, as a way of studying the main processes without the less important or poorly understood processes that cannot be removed from an experiment.

1D is best. The channels of a monolith are clearly intended to operate as small parallel reactors, and ideally identically. If they did a poor job of this, we would have devised a fix by now to stop the waste of precious metal. Small deviations from this idealization that are within the accuracy constraints of the chemistry are not worth the effort. The better use of 3D is to design the adjoining components (insulating can, cones, urea injection system, etc.) to ensure approximate 1D behavior of the key ATR component. An ATR that is not approximately 1D is a poor performer, so apportion your attention and efforts accordingly. Even when studying nonuniform ammonia concentration, the main effect is through a trace species of the mixture, so alternatives short of a full 3D flow solution for the ATR are possible.

Don't squander degrees of freedom. When generating or calibrating global kinetics, many pitfalls center on expecting more than needed. Global rates can only capture part of the "truth", and we don't want them any more complex than needed. Make sure you only constrain global rates to agree with realistic results. For example, optimization software can easily generate unrealistic concentration combinations such as too far from stoichiometric, or very high CO with very low HCs. Sometimes temperature sweeps

are too broad, thereby capturing rates that are too small at low temperature to matter in practice, or too high at high temperature to survive competition with transport limitations. These tests are fine for academic work, but waste valuable degrees of freedom better used in the practical regime.

Enhancing mass transfer kills you for momentum. Proposals spring up regularly to improve mass transfer by introducing various flow obstructions to perturb the purely axial flow of monolith channels. However, these geometrical changes also necessarily “enhance” momentum transfer proportionally, which manifests as increased shears and higher pressure drops. Since the increased concentration gradients and increased shears are manifestations of the same process for species and momentum, respectively, they cannot be separated. The extra pressure drop from the work of forcing the flow around these obstacles is no more acceptable in your production system than it was in any of the previous such proposals.

Beware of ATRs as on/off switches. ATRs are not black boxes, but integral reactors whose conversion of reactants increases with axial distance. However, if the conversion is either 0% or 100%, we can only bound the reaction rate. The quantification we need for kinetics work is lost if we lack partial conversion, so choose your experimental conditions carefully to ensure this. Similarly, it is unfortunately too common to see model predictions of cumulative emissions over a drive cycle in which the reactor is nearly always at 0% or 100% conversion. This is a very poor test of the kinetics. Always plot inlet concentrations on the same graph so you can spot this.

Use academic work judiciously. The practical ATR modeler is reliant on the scholarship of academic research. Sometimes this work has involved studying effects that we know are present, but were not yet known to be essential in our modeling (for example, temperature gradients across a filter wall or concentration gradients across a washcoat layer). When modeling practical applications, we should appreciate that the academic researcher, having devoted substantial resources to investigating a potential effect, has no particular incentive to emphasize a null result. You might have to dig for it. My example of filter wall temperature gradients is actually a counterexample, because several old but important papers have indeed emphasized the null result, to the subsequent benefit of the aftertreatment modeling community. My other example of washcoat concentration gradients is more current and less straightforward, largely due to the difficulty of estimating effective diffusion resistance within a washcoat. An interesting case study is to compare two often-referenced papers from the same research group (Hayes, et al., *Applied Catalysis B: Environmental* 25 (2000) 93-104 and Zhang et al., *Trans IChemE, Part A, Chemical Engineering Research and Design*, 2004, 82(A4) 481-489). In the second paper, the authors attribute an order of magnitude difference in diffusion coefficient to the more typical presence of macropores in their sample which were absent in the sample from the first paper. Various subsequent literature studies have quoted the higher diffusion resistance from the first paper to establish a limit on the range and then proceeded to generate results based on this limiting value. The practical modeler must also know the significance of the second paper and that macropores are expected in their intended washcoats. Therefore conclusions regarding pore diffusion resistance based only on estimates from the first paper overestimate the effect by an order of magnitude for such intended applications.

More is not better. Try to study speculative secondary processes separately, outside of your mainstream ATR modeling. If you include everything you can think of at all times, it bogs down your calculation and confuses interpretation of your results.