



Accelerating simulation and optimisation of cyclic adsorption processes with differentiable programming

Alex Glover¹, Maria M. Papathanasiou^{1*}, Ronny Pini^{1**}

¹Department of Chemical Engineering, Imperial College London, South Kensington Campus, London, SW7 2AZ, United Kingdom

The Sargent Centre for Process Systems Engineering, Imperial College London, South Kensington Campus, London, SW7 2AZ, United Kingdom

ARTICLE INFO

Keywords:

Differentiable programming
Automatic differentiation
Cyclic processes
Pressure vacuum swing adsorption
Multi-objective optimisation

ABSTRACT

The design of cyclic adsorption processes is computationally demanding, requiring repeated convergence to cyclic steady state within an iterative optimisation loop. Conventional workflows treat the process simulator as a black box and rely on derivative-free optimisation, resulting in design campaigns that can require hundreds to thousands of CPU hours. This work presents an end-to-end differentiable model of a pressure vacuum swing adsorption process, developed using the JAX differentiable programming framework and applied here to a benchmark post-combustion carbon capture problem. Automatic differentiation provides exact gradients throughout the entire computational workflow. The differentiation of a single process cycle provides the Jacobian for a Newton iteration to decrease both the number of iterations and the simulation time required to reach cyclic steady state by a factor of 20 relative to a representative MATLAB implementation. Exact gradients of the performance metrics with respect to the design variables further enable gradient-based multi-objective optimisation using the IPOPT algorithm. Applied to a six-variable design problem, the latter produces a superior Pareto front with improved coverage of the trade-off space and closer convergence to the optimal front than the genetic algorithm NSGA-II. Notably, the full front is obtained over two orders of magnitude faster than the conventional approach. By retaining the full mechanistic model while making it differentiable, this framework transforms cyclic adsorption process design from slow black-box simulation with derivative-free optimisation to efficient gradient-enhanced modelling and optimisation, enabling rapid and systematic exploration of complex design spaces.

1. Introduction

Adsorption-based processes are widely employed in diverse gas separation applications and have received increasing interest for post-combustion and other carbon dioxide capture technologies (Raganati et al., 2021; Pini et al., 2026). The process itself is operated through a sequence of steps, in which either the pressure or the temperature, or both, are varied to selectively adsorb and desorb the components in the gas mixture. Mechanistic adsorption models, such as the one-dimensional axially dispersed plug flow model, have been used extensively throughout the literature for experiment design and interpretation, and to describe the behaviour observed in cyclic processes (Ward and Pini, 2022b; Krishnamurthy et al., 2014; Todd et al., 2001). These models have been used for open-loop predictions to screen adsorbents (Rajagopalan et al., 2016) and for integrated adsorbent-process optimisation (Khurana and Farooq, 2017). The design exercise is not trivial and requires finding the set of operating parameters (e.g., the

duration and operating pressure of each step in the cycle) that balance conflicting objectives (or key performance indicators, KPIs) such as the purity and recovery of the product as well as the productivity and energy consumption of the process. This design workflow is typically formulated as a multi-objective optimisation problem that incorporates a combination of these four KPIs, or as a single-objective economic optimisation (Ward and Pini, 2022a; Haghpanah et al., 2013). Recent years have seen developments that greatly expand the design space which needs to be explored, including novel adsorbent materials (e.g., metal organic frameworks (Mahajan and Lahtinen, 2022)), multi-sorbent beds (Ward and Pini, 2024), structured contactors (Antunes Morgado and Mazzotti, 2026) as well as process configurations featuring additional steps beyond the foundational 2-bed, 4-step Skarstrom cycle (Pai et al., 2025). The optimisation of these more advanced designs increases the computational demand, with individual problems often requiring hundreds to thousands of CPU hours (Ward and Pini, 2022a; Peh et al., 2022; Wilkins et al., 2025; Subraveti et al., 2019; Pai et al., 2020a).

* Corresponding author.

** Corresponding author at: Department of Chemical Engineering, Imperial College London, South Kensington Campus, London, SW7 2AZ, United Kingdom.
E-mail addresses: maria.papathanasiou11@imperial.ac.uk (M.M. Papathanasiou), r.pini@imperial.ac.uk (R. Pini).

<https://doi.org/10.1016/j.compchemeng.2026.109904>

Received 13 July 2026; Received in revised form 14 September 2026; Accepted 16 September 2026

Available online 22 September 2026

0098-1354/© 2026 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

The ability to accurately determine model gradients underpins several applications of the mechanistic adsorption model just described — from sensitivity analysis to optimisation. This need becomes critical as the total number of parameters increases and when dealing with highly non-linear processes, as it is the case for the design of adsorption-based separations. Differentiable programming, a framework in which the numerical programme is constructed such that it can be *automatically* differentiated throughout (i.e. without manual derivation), offers a route to drastically increase the computational efficiency of these tasks. The benefits of differential programming have been demonstrated in several application areas, including crystallisation (Alsubeihi et al., 2025), computational fluid dynamics (Bezgin et al., 2023), chemical kinetics (Lee et al., 2025), and molecular dynamics (Schoenholz and Cubuk, 2021). In this work, we develop a differentiable framework for the simulation and optimisation of cyclic adsorption processes. A one-dimensional axially dispersed plug flow model is implemented in JAX, leveraging just-in-time compilation, automatic differentiation, and accelerated linear algebra throughout the process model. These capabilities are used to accelerate the numerical solution, the convergence to cyclic steady state, and the multi-objective optimisation. The following subsections review the modelling and optimisation approaches used within the adsorption community and contrast these with methods from the process systems engineering literature, highlighting the gap that differentiable programming can address.

1.1. Adsorption process design

The physical phenomena in the fixed-bed adsorption column are described by mass and energy balances alongside equations characterising the local adsorption behaviour, yielding a system of partial differential equations (PDEs). A method of lines approach, typically using a finite volume scheme, is often used to discretise the system of PDEs into a set of ordinary differential equations (ODEs) (Haghpanah et al., 2013). Upon application of the relevant boundary conditions, these equations are solved for each step of the adsorption process using an implicit solver due to the stiffness of the system (LeVeque, 2002). Owing to the repeated sequence of steps imposed on the adsorption column, the process is designed to operate at cyclic steady state (CSS) — an operating point for which the state of the column at any given time matches its state at the exact same time in the previous cycle. Reaching CSS typically involves simulation from a startup condition (a clean or saturated bed) through repeated cycles until periodic stability is achieved. Adsorption processes with short cycle times and pressure rather than temperature swings may require hundreds of cycles to reach CSS (Croft and LeVan, 1994a).

The optimisation of adsorption processes has predominantly been carried out using derivative-free or black box methods due to their simplicity and their ability to deal with non-convex optimisation problems. Here, the cyclic process model is iterated until CSS is reached; the objective function and performance constraints (typically the KPIs) are then evaluated, and the optimisation algorithm computes the next set of parameters for evaluation. This structure is displayed in Fig. 1a. A genetic algorithm, such as the non-dominated sorting genetic algorithm II (NSGA-II) (Deb et al., 2000), is by far the most widely used derivative-free approach in the adsorption literature (Haghpanah et al., 2013; Wilkins et al., 2025; Subraveti et al., 2019; Ward and Pini, 2024; Pai et al., 2020b; Beck, 2014; Liao et al., 2025). In practice, this approach requires simulating thousands of candidate designs, each of which must be iterated to CSS, resulting in the large computational times reported in the literature (Ward and Pini, 2022a; Peh et al., 2022; Wilkins et al., 2025; Subraveti et al., 2019; Pai et al., 2020a). Other derivative-free approaches, such as Bayesian optimisation, have been also applied, demonstrating reduced computational times relative to genetic algorithms (Hao et al., 2021; Ward and Pini, 2022a). Alternatively, many studies use simplified models (Khurana and Farooq, 2016; Subramanian Balashankar et al., 2019) or machine learning surrogates

trained from a set of high-fidelity simulations (Subraveti et al., 2019; Chatziasteriou et al., 2025) to reduce the computational cost which are often formulated in a way that alleviates the requirement of iterating to CSS.

1.1.1. Gradient-based methods for cyclic process optimisation

The process systems engineering literature reports on a range of gradient-based methods for the optimisation of periodic processes that address many of the computational limitations described above, though these have seen limited adoption in adsorption applications. Early approaches accelerated convergence to CSS using Newton or quasi-Newton iterations applied to the direct determination of the periodic state (Croft and LeVan, 1994a; Ding and LeVan, 2001). These methods achieve rapid (quadratic) convergence, but require evaluation of the Jacobian matrix to describe the sensitivities of the final state of the cycle with respect to the initial state. While convergence to CSS is achieved in much fewer iterations than successive substitution, the additional computational burden of the Jacobian evaluation has limited their adoption (Ko et al., 2003; Ding and LeVan, 2001; Jiang et al., 2003). Jacobian-free vector-based extrapolation methods have been proposed (Subraveti et al., 2024, 2026), achieving up to fourfold speed-ups while remaining simple to implement. However, their convergence rate remains inferior to exact Newton iteration methods. Machine learning models have also been used to predict approximate CSS profiles, providing improved initial guesses that accelerate convergence in subsequent simulations (Pai et al., 2020a).

An alternative strategy for the optimisation of periodic processes is full discretisation of both space and time (Nilchan and Pantelides, 1998), transforming the governing PDEs into a nonlinear algebraic system of equations solved via nonlinear programming. While it enables direct inclusion of the CSS criterion and the optimisation objectives (Fig. 1b), this approach is highly sensitive to the initial guess and can suffer from convergence issues for complex process models (Ko et al., 2003; Jiang et al., 2003). A pseudo-transient formulation has been proposed to mitigate this by gradually introducing system dynamics, improving robustness (Tsay et al., 2018).

In practice, the single discretisation approach — where PDEs are discretised only in space and the resulting ODEs are integrated in time — is more commonly used. Within this framework, simultaneous optimisation formulations incorporate the CSS convergence criterion as an inequality constraint (Fig. 1c) (Jiang et al., 2003). Although it avoids explicit iteration to CSS, this approach increases problem size and complexity, as both the state variables at every discretised point and the design variables become decision variables. These formulations rely on gradient-based algorithms, requiring sensitivity information that can be obtained via finite differences (Cruz et al., 2005), adjoint ODEs (Latifi et al., 2008), or through discrete methods (Jiang et al., 2003). Many of these optimisation problems have been solved using the reduced Hessian Sequential Quadratic Programming (SQP) algorithm (Ternet and Biegler, 1998; Biegler et al., 2005). Other works have formulated the design problem as an optimal control problem using dynamic optimisation approaches (Agarwal et al., 2010; Dowling et al., 2012; Khajuria and Pistikopoulos, 2013).

The computation of sensitivities is typically the dominant computational cost in the gradient-based approaches described above, often exceeding 90% of the CPU time (Jiang et al., 2003; Kawajiri and Biegler, 2006), and motivating the development of optimisation algorithms that can accommodate inexact Jacobians (Vetukuri et al., 2010). Despite these advances, the need for accurate sensitivities — coupled with approximation errors, stability concerns, or implementation effort — has hindered broader adoption. Automatic differentiation within a differentiable programming framework offers exact, stable sensitivities without manual derivation. While differentiable programming has been successfully applied in adjacent process engineering domains (Alsubeihi et al., 2025; Bezgin et al., 2023; Lee et al., 2025; Schoenholz and Cubuk, 2021), its application to the rigorous simulation and optimisation of cyclic adsorption processes has not been demonstrated.




Fig. 1. Optimisation structures used within the literature: derivative-free optimisation (a), complete discretisation (b), and simultaneous tailored approach (c). Source: Adapted from Biegler et al. (2005).

2. A differentiable framework for adsorption processes

In this section, we introduce the differentiable framework for the simulation and optimisation of cyclic adsorption processes. The framework follows the same modelling approach used throughout the adsorption literature: the governing PDEs are discretised with a finite volume scheme, the resulting ODEs are integrated with an implicit solver for each step of the cycle, and the process is iterated to CSS from which the KPIs are computed for use within a multi-objective optimisation problem. The mechanistic model is therefore unchanged. The new development is in making the entire computational workflow differentiable, enabling the use of gradient information for the accelerated simulation and optimisation of the adsorption process.

As identified in the preceding section, the conventional model formulation in the adsorption literature does not exploit gradient information at any stage of the design process. The differentiable framework addresses this at three levels, as illustrated in Fig. 2. First, at the level of the ODE solution, just-in-time (JIT) compilation and automatic differentiation (AD) within the implicit solver are used to accelerate the numerical integration and provide sensitivities of the solution with respect to the initial state and model parameters (green dashed arrows in the figure). Second, these sensitivities enable a Newton-based direct determination method for convergence to CSS, replacing the successive substitution approach (dashed green box in the figure). Third, the CSS condition is treated as an implicit function, such that the model outputs both objective functions (e.g., KPIs) and their gradients with respect to the decision variables for use within gradient-based multi-objective optimisation. These three components are described in the following subsections.

2.1. Solution and sensitivities

The system of PDEs that describes the adsorption process (Table 1) is discretised via a finite volume scheme where the fluxes at the walls are evaluated using a third-order weighted essentially non-oscillatory method (Haghpour et al., 2013). The resulting dimensionless discretised equations are given in the Section S2 of the Supporting Information. Eq. (1) describes their general form, where $\mathbf{x}$ denotes the vector

of dimensionless state variables (pressure, temperature, mole fractions, column wall temperature and adsorbed quantities) at each volume node along the column, θ denotes the model parameters, and $\mathbf{x}_0$ is the state of the column at the start of the step.

$$\frac{d\mathbf{x}}{dt} = f(t, \mathbf{x}(t), \theta), \quad \mathbf{x}(0) = \mathbf{x}_0 \quad (1)$$

Fig. 3 illustrates the flowchart for the solution of one process cycle as well as the sensitivities computation and propagation through the model solution. An implicit ODE solver is chosen for the integration of the states due to the stiffness of the equations resulting from the discretised second order spatial derivatives (diffusion and conduction) and the adsorption and heat transfer source/sink terms (LeVeque, 2002). At each time step of the implicit solver, a system of non-linear equations must be solved, typically via Newton iterations, which in turn requires the evaluation of the Jacobian of the right-hand side function $f(t, \mathbf{x}(t), \theta)$ with respect to the state vector $\mathbf{x}$. With the differentiable framework adopted here, this Jacobian is computed using AD rather than being approximated with finite differences or derived manually.

A key requirement for the gradient-based methods used in this work is the ability to obtain sensitivities of the solution with respect to the initial state ($\partial \mathbf{x}_k / \partial \mathbf{x}_k^0$) and model parameters ($\partial \mathbf{x}_k / \partial \theta$) – see Fig. 3. For implicit solvers, such as the one used in this work, a naive use of AD would differentiate through all Newton iterations at each time step, leading to high computational cost and numerical sensitivity to the solver convergence tolerance. Instead, the sensitivities are computed at convergence of the non-linear system as illustrated in Fig. 3. This task is achieved by directly differentiating the numerical time-stepping scheme using the discrete sensitivity method, implemented in the forward direction (Sapienza et al., 2024). The resulting tangent gradients are then “manually” propagated alongside the state variables throughout the simulation. To enable comparison with conventional approaches used in adsorption modelling, we have implemented in JAX a version of the numerical differentiation formula (NDF) method used in MATLAB’s ode15s solver (Shampine and Reichelt, 1997). The corresponding state and sensitivity equations are given in Section S1 of the Supporting Information. Notably, in contrast to ode15s, the Jacobian is re-evaluated at every accepted step to ensure accurate




Fig. 2. Differentiable adsorption model framework. The sensitivity propagation for the computation of the gradients of the objective functions (e.g., KPIs) and the constraints with respect to the decision variables are shown with the green arrows. The two options for the iteration to CSS are illustrated, namely successive substitution (dashed black box) and direct determination (dashed green box).

sensitivity calculations, rather than being reused at many steps when sensitivities are not required.

The Jacobian matrices required in both the state and sensitivity equations exhibit a known sparsity pattern arising from the spatial structure of the discretised transport equations for the adsorption column. The sparsity is exploited by partitioning the right-hand side function $f(t, \mathbf{x}(t), \theta)$ so that each evaluation operates only on the subset of volumes that directly influence a given spatial location. AD is then applied locally to these reduced functions, computing only the non-zero derivative entries. These operations are vectorised across all control volumes using JAX's `vmap` function. The linear systems arising in both the Newton iterations and the sensitivity equations share the same sparsity pattern and are solved using a sparse linear solver (KLU algorithm (GDSFactory, 2026)). This exploitation of sparsity ensures that the approach scales to finer spatial discretisation.

2.2. Direct determination of cyclic steady state

Adsorption processes are typically cyclic in nature where each cycle is composed of a sequence of steps for which the end state of a cycle becomes the initial state of the following cycle. In modelling studies, a mathematical criterion is used to define whether a process is operating at CSS. Throughout the literature, authors have enforced this criterion based upon change in the values of the state variables (Nilchan and Pantelides, 1998), KPIs (Marcinek et al., 2020) and the overall mass

balance for a cycle (Krishnamurthy et al., 2014). In this work we define the cycle error as the difference between the vector of state variables at the start and end of a cycle:

$$\mathbf{e}_k = \mathbf{x}_k - \mathbf{x}_k^0 \quad (2)$$

The CSS criterion is satisfied when all elements of this vector, normalised to the initial state $\mathbf{x}_k^0$, are below a predefined tolerance (for example, 10^{-5}). Typically, the so-called successive substitution method is used to reach the CSS. This method involves simulating the process from an initial condition (a bed saturated with the light or heavy component) to CSS. Mathematically, this is known as a Picard iteration where the output of the cycle function is used as the input to the next cycle. This method has the benefit of following the dynamics of the process from startup to CSS – information not necessarily required for optimisation studies where the target is a design at CSS. This method can be slow, especially for systems where the cycle time is short, which is common in pressure swing processes. The slow convergence is described by the CSS having a spectral radius, the largest eigenvalue of the cycle Jacobian matrix $\partial \mathbf{x}_k / \partial \mathbf{x}_k^0$ and also known as the dominant Floquet multiplier, close to, but less than 1 (Douglas LeVan and Croft, 1995). In contrast, the direct determination method does not follow the process dynamics; rather it predicts the input to the next cycle from a Newton iteration step (Croft and LeVan, 1994a):

$$\mathbf{x}_{k+1}^0 = \mathbf{x}_k^0 - \alpha \frac{\partial \mathbf{e}_k}{\partial \mathbf{x}_k^0}^{-1} \mathbf{e}_k \quad (3)$$




Fig. 3. Flowchart for the state and the sensitivities computation during the solution of one adsorption cycle.

where the error, e_k , and the error gradient with respect to the initial state, $\partial e_k / \partial x_k^0$, are computed at each iteration k . Fig. 2 illustrates the flowchart for the differentiable model, where the iteration to CSS is shown with the loop. The traditional successive substitution and the direct determination methods are shown. The error gradient is computed as follows:

$$\frac{\partial e_k}{\partial x_k^0} = \frac{\partial x_k}{\partial x_k^0} - I \quad (4)$$

Again, the differentiable framework described above allows for the efficient evaluation of the full Jacobian matrix for the sensitivities of the states at the end of the cycle with respect to the initial states at each iteration. As the Newton iteration does not follow the process dynamics, it can generate initial conditions for the model that include non-physical values (Croft and LeVan, 1994a), such as negative mole fractions or adsorbed quantities. To prevent failure of the ODE solver, a damped Newton iteration is implemented in Eq. (3), which uses a

damping factor $\alpha \in [0, 1]$ to scale the Newton step. This amounts to a line search along the direction of the Newton step, in which the step length is chosen so that the next iterate remains physically admissible: the search identifies the fraction of the full Newton step at which the first state variable would reach zero, and retreats from it. The damping factor is:

$$\alpha = \frac{1}{2} \min_{i \in \mathcal{N}_k} \left(\frac{x_{k,i}^0}{x_{k,i}^0 - x_{k+1,i}^{0,u}} \right) \quad (5)$$

where $x_{k+1,i}^{0,u}$ is the i th component of the state vector for the next iteration obtained from the undamped Newton iteration ($\alpha = 1$), and $x_{k,i}^0$ is the corresponding component of the current iterate. The set $\mathcal{N}_k = \{i : x_{k+1,i}^{0,u} < 0\}$ collects the indices of those discretised state variables that the undamped step would give as negative, such as mole fractions or adsorbed quantities. The minimum over $\mathcal{N}_k$ selects the first state variable to reach zero along the Newton direction, and the factor

of $\frac{1}{2}$ provides a safety margin so that the next iterate is bounded away from zero rather than merely non-negative.

2.3. Gradient-based optimisation

The gradient-based optimisation methodology uses the same structure as the derivative-free optimisation studies illustrated in Fig. 1a, where the iteration to CSS is the inner loop within each evaluation. Without losing generality and with reference to Fig. 2, Eq. (6) describes an optimisation problem involving a single objective function ϕ , alongside a set of inequality constraints $\mathbf{g}$, which are both typically described through the evaluation of the KPIs of the process. The decision variables for the optimisation θ are bounded (the subscript L and U refer to the lower and upper bound, respectively), and $\mathbf{x}^*$ is the vector of the state variables at the start of the cycle where the CSS criterion is satisfied.

$$\begin{aligned} \min_{\theta} \quad & \phi(\mathbf{x}^*, \theta) \\ \text{s.t.} \quad & \mathbf{g}(\mathbf{x}^*, \theta) \leq 0 \\ & \theta \in [\theta^L, \theta^U] \end{aligned} \quad (6)$$

The CSS condition is expressed as an implicit function:

$$\mathbf{x}^* - \mathbf{f}_{\text{cycle}}(\mathbf{x}^*, \theta) = 0 \quad (7)$$

where the function $\mathbf{f}_{\text{cycle}}$ describes the whole process cycle. At each iteration, the gradients of the objective function and the constraints with respect to the design parameters are computed as follows:

$$\frac{d\phi}{d\theta} = \frac{\partial\phi}{\partial\theta} + \frac{\partial\phi}{\partial\mathbf{x}^*} \frac{d\mathbf{x}^*}{d\theta} \quad (8)$$

$$\frac{d\mathbf{g}}{d\theta} = \frac{\partial\mathbf{g}}{\partial\theta} + \frac{\partial\mathbf{g}}{\partial\mathbf{x}^*} \frac{d\mathbf{x}^*}{d\theta} \quad (9)$$

where $d\mathbf{x}^*/d\theta$ is evaluated through the implicit function theorem:

$$\frac{d\mathbf{x}^*}{d\theta} = \left(\mathbf{I} - \frac{\partial\mathbf{f}_{\text{cycle}}}{\partial\mathbf{x}^*} \right)^{-1} \frac{\partial\mathbf{f}_{\text{cycle}}}{\partial\theta} \quad (10)$$

Fig. 2 illustrates how these gradients are applied in the differentiable model framework developed in this work. We note that the implicit nature of the CSS criterion implies that the optimiser does not ‘see’ this constraint. As a result, any error associated with the CSS tolerance can lead to noise in the gradients. In this work, a CSS constraint tolerance of 10^{-5} has been used throughout. However, for systems where the spectral radius approaches 1 and tighter optimality tolerances are used within the optimiser, a tighter CSS tolerance might be required.

3. Pressure vacuum swing adsorption case study

The case study considered in this work is a four-step pressure vacuum swing adsorption (PVSA) process for the separation of CO_2 from a dry flue gas containing 15 mol% CO_2 in N_2 (number of species, $n_c = 2$) at atmospheric pressure and 298 K, using Zeolite 13X as the adsorbent. This system has been the subject of numerous optimisation studies owing to its relevance to post-combustion carbon capture and is therefore a good benchmark for this work (Haghpahan et al., 2013; Ward and Pini, 2022a; Subraveti et al., 2019). The four-step cycle consists of (Fig. 4): adsorption at high pressure during which the feed is introduced at one end of the column and the light product (N_2) exits at the other; co-current blowdown to an intermediate pressure which removes most of the residual N_2 from the column; reverse evacuation to the low vacuum pressure during which the heavy product (CO_2) is collected; and feed pressurisation back to the adsorption pressure.

A one-dimensional, non-isothermal, non-isobaric model describes the coupled heat and mass transfer within the packed adsorption column. The model comprises partial differential equations for the component mass balances, the overall mass balance, the column internal energy balance, the column wall energy balance and the steady-state momentum balance in the form of Darcy’s law for the pressure drop along the column. The following assumptions are adopted in deriving the governing equations:




Fig. 4. Four-step PVSA cycle. Each step from left to right is: Adsorption (ads), Blowdown (bd), Reverse Evacuation (evac), Feed Pressurisation (pres). In our case study, the feed is a mixture of CO_2 and N_2 , the light product is the N_2 -rich stream and the heavy product is the CO_2 -rich stream. P_H , P_1 and P_L refer to the high, intermediate and low pressure. The duration of each step is indicated by the horizontal arrows.

- Only axial gradients are considered (radial gradients within the bed are neglected). The suitability of this formulation has been demonstrated through both laboratory (Marx et al., 2016; Estupiñán Perez et al., 2019) and pilot-scale studies (Krishnamurthy et al., 2014; Henrotin et al., 2026).
- The gas phase obeys the ideal gas law.
- Local thermal equilibrium exists between the gas, adsorbed and solid phases within the bed.
- The heat capacities of the gas, adsorbed and solid phases are constant.
- The densities of the adsorbent and the column wall are constant.
- The column void fraction and adsorbent particle radius are uniform and constant.
- The axial mass dispersion coefficient and effective thermal conductivity are constant.
- The pressure drop along the column is described by Darcy’s law.

The governing equations are summarised in Table 1. The state variables are the total pressure p , the gas-phase mole fractions y_i of each component i apart from the last, the gas temperature T , the column wall temperature T_w , and the solid-phase loadings q_i for all components. In this work, the PDEs are discretised with 10 finite volumes. These equations are accompanied by the adsorption isotherms which describe the thermodynamic equilibrium between the gas and the adsorbed phase. For the Zeolite 13X material used in this work, the extended dual-site Langmuir model is used to quantify this relationship for the CO_2/N_2 mixture:

$$q_i^* = \frac{q_{b,i} b_i c_i}{1 + \sum_i^{n_c} b_i c_i} + \frac{q_{d,i} d_i c_i}{1 + \sum_i^{n_c} d_i c_i} \quad (11)$$

where $q_{b,i}$ and $q_{d,i}$ are the saturation capacities at each adsorption site for each component, $c_i = y_i p / RT$ is the concentration of each component in the gas phase, and b_i and d_i are the equilibrium constants:

$$b_i = b_{0,i} e^{-\frac{\Delta U_{b,i}}{RT}} \quad (12)$$

Table 1

Mass, energy and momentum balances for the 1-dimensional axially dispersed plug flow model.

Overall material balance
$\frac{\partial p}{\partial t} - \frac{p}{T} \frac{\partial T}{\partial t} = -T \frac{\partial}{\partial z} \left(\frac{pv}{T} \right) - RT \frac{\rho_b}{\epsilon} \sum_{i=1}^{n_c} \frac{\partial q_i}{\partial t}$
Component material balance for species $i \in \{1, \dots, n_c - 1\}$
$\frac{\partial y_i}{\partial t} = \frac{T}{p} D_L \frac{\partial}{\partial z} \left(\frac{p}{T} \frac{\partial y_i}{\partial z} \right) - \frac{T}{p} \left(\frac{\partial}{\partial z} \left(\frac{y_i pv}{T} \right) - y_i \frac{\partial}{\partial z} \left(\frac{pv}{T} \right) \right) - \frac{RT}{p} \frac{\rho_b}{\epsilon} \left(y_i \sum_{i=1}^{n_c} \frac{\partial q_i}{\partial t} - \frac{\partial q_i}{\partial t} \right)$
Solid phase material balance for species $i \in \{1, \dots, n_c\}$
$\frac{\partial q_i}{\partial t} = k(q_i^* - q_i)$
Column energy balance
$\left[C_{p,g} \frac{P}{RT} + \frac{\rho_b}{\epsilon} (C_{p,s} + C_{p,a} \sum_{i=1}^{n_c} q_i) \right] \frac{\partial T}{\partial t} = \frac{K_z}{\epsilon} \frac{\partial^2 T}{\partial z^2} - \frac{C_{p,g}}{R} \left(\frac{\partial(pv)}{\partial z} - T \frac{\partial}{\partial z} \left(\frac{pv}{T} \right) \right) - \frac{\rho_b}{\epsilon} T (C_{p,a} - C_{p,g}) \sum_{i=1}^{n_c} \frac{\partial q_i}{\partial t} + \frac{\rho_b}{\epsilon} \sum_{i=1}^{n_c} \left(-\Delta H_i \frac{\partial q_i}{\partial t} \right) - \frac{2h_{in}}{\epsilon r_{in}} (T - T_w)$
Column wall energy balance
$\rho_w C_{p,w} \frac{\partial T_w}{\partial t} = K_w \frac{\partial^2 T_w}{\partial z^2} + \frac{2r_{in} h_{in}}{r_{out}^2 - r_{in}^2} (T - T_w) - \frac{2r_{out} h_{out}}{r_{out}^2 - r_{in}^2} (T_w - T_a)$
Steady-state momentum balance (Ergun equation)
$-\frac{\partial p}{\partial z} = \frac{150}{4r_p^2} \left(\frac{1-\epsilon}{\epsilon} \right)^2 \mu v$

$$d_i = d_{0,i} e^{-\frac{\Delta U_{a,i}}{RT}} \quad (13)$$

The isotherm parameters are taken from the unequal energy sites fitting from [Wilkins and Rajendran \(2019\)](#). The mass transfer resistance for the diffusion from the bulk gas to the adsorption site is described through the linear driving force (solid phase material balance in [Table 1](#)) with a constant mass transfer coefficient, k . Lastly, the longitudinal dispersion coefficient is found through the following correlation ([Ruthven, 1984](#)).

$$D_L = 0.7D_m + v_{feed} r_p \quad (14)$$

where D_m is a molecular diffusion coefficient and v_{feed} is the interstitial feed velocity. Values of the parameters used within this model are given within [Table 2](#). The boundary conditions for each step are set according to the direction of flow and the type of operation. During the adsorption step, the feed-end boundary imposes Danckwerts conditions on the mole fractions and temperature, and a fixed inlet velocity. During blowdown and evacuation steps, an exponential pressure profile is imposed at the product end of the column, whilst zero-flux conditions are applied at the closed end. During the pressurisation step, an exponential pressure profile is prescribed at the feed boundary. The boundary conditions for each step are described in Section S5 of the Supporting Information.

3.1. Isotheric heat of adsorption

The isosteric heat of adsorption describes the enthalpy released during adsorption. The Clausius-Clapeyron equation for the isosteric heat of adsorption for a multicomponent ideal gas mixture in equilibrium with an adsorbent is given as ([Sircar, 1992](#))

$$-\Delta H_{iso,i} = RT^2 \left(\frac{d \ln(P y_i)}{dT} \right)_q = RT^2 \left[\frac{1}{P} \left(\frac{dP}{dT} \right)_q + \frac{1}{y_i} \left(\frac{dy_i}{dT} \right)_q \right] \quad (15)$$

The subscript $q = (q_1, \dots, q_{n_c})$ indicates that the derivative is evaluated at constant adsorbed-phase loading of every component. As the DSL isotherm describes loading as a function of the gas phase pressure, temperature and composition, we apply implicit differentiation to this isotherm using AD in JAX. At constant loading ($dq_i = 0$), the following expression is derived for all components:

$$0 = \frac{\partial q_i}{\partial P} \frac{dP}{dT} + \frac{\partial q_i}{\partial T} + \sum_{j=1}^{n_c-1} \frac{\partial q_i}{\partial y_j} \frac{dy_j}{dT}, \quad (16)$$

Eq. (16) is solved simultaneously for each component $i = 1..n_c$, whereby $(dy_{n_c}/dT)_q = -\sum_{j=1}^{n_c-1} (dy_j/dT)_q$. The derivatives of pressure, and all mole fractions with respect to temperature at constant loading are used in Eq. (15) to estimate the isosteric heat of adsorption.

3.2. Key performance indicators

The four KPIs that are used for the design of the PVSA process include the purity of the CO₂ product (Pu_{CO_2}), the recovery of CO₂ (Re_{CO_2}), the productivity of the process (Pr) and the energy consumption associated with the compressors and vacuum pumps (E_T). For the calculation of these KPIs, three further process units are integrated in time along with the states of the column. These include an isentropic compressor, an isentropic vacuum pump and a tank to calculate the quantity and composition of gas discharged from the outlet of the adsorption step. The power consumption of the compressor and vacuum pump as well as the total gas and CO₂ molar flowrate from the outlet of the adsorption step are added as state variables for the ODE solution. The equations for these units and the KPIs are given in Sections S7 and S8 of the Supporting Information.

3.3. Process optimisation

The design of adsorption processes for post combustion carbon capture applications typically involve multi-objective optimisation of two conflicting KPIs. The first involves the maximisation of the purity and recovery, which we formulate as a constrained single-objective optimisation problem:

$$\begin{aligned} \min_{\theta} \quad & -Re_{CO_2}(\theta) \\ \text{s.t.} \quad & Pu_{CO_2}(\theta) \geq \epsilon \\ & \theta \in [\theta^L, \theta^U] \end{aligned} \quad (17)$$

where ϵ is suitably varied so as to identify the Pareto curve of optimal solutions (epsilon-constraint method ([Miettinen, 1998](#))). The second optimisation problem uses an analogous formulation to maximise productivity and minimise energy consumption, albeit under constraints on purity and recovery:

$$\begin{aligned} \min_{\theta} \quad & E_T(\theta) \\ \text{s.t.} \quad & Pr(\theta) \geq \epsilon \\ & Pu_{CO_2}(\theta) \geq 95\% \\ & Re_{CO_2}(\theta) \geq 90\% \\ & \theta \in [\theta^L, \theta^U] \end{aligned} \quad (18)$$

For these optimisation problems, the design variables are bounded with the values given in [Table 3](#), and include the duration and pressure of the adsorption step (t_{ads} and P_H , respectively), the feed velocity (v_{feed}), the duration and pressure of the blowdown (t_{bd} and P_L , respectively), and the duration of the evacuation step (t_{evac}). For the design

Table 2
Parameters used within adsorption model simulation.

Parameter	Units	Value
Column length, L	[m]	1.0
Column inner radius, r_{in}	[m]	0.145
Column outer radius, r_{out}	[m]	0.162
Bed density, ρ_b	[kg/m ³]	712
Bed voidage, ϵ		0.37
Particle voidage, ϵ_p		0.35
Particle radius, r_p	[m]	0.001
Particle tortuosity, τ_p		3
Molecular diffusivity, D_m	[m ² /s]	1.5×10^{-7}
Thermal conductivity of gas, K_z	[J/m/K/s]	0.09
Thermal conductivity of wall, K_w	[J/m/K/s]	16
Heat capacity of gas phase, C_{pg}	[J/mol/K]	30.7
Heat capacity of adsorbed phase, C_{pa}	[J/mol/K]	30.7
Heat capacity of adsorbent, C_{ps}	[J/kg/K]	1070
Heat capacity of column wall, C_{pw}	[J/kg/K]	502
Density of wall, ρ_w	[kg/m ³]	7800
Dynamic viscosity of gas, μ	[kg/m/s]	1.72×10^{-5}
Overall inside heat transfer coefficient, h_{in}	[J/m ² /K/s]	8.6
Overall outside heat transfer coefficient, h_{out}	[J/m ² /K/s]	2.5
Pressure profile time constant, λ	[s]	0.5
Compressor efficiency, η		0.72
Mass transfer coefficient (LDF), k	[s]	26.25
Saturation capacity of CO ₂ at site 1, q_{b,CO_2}	[mol/kg]	3.257
Saturation capacity of CO ₂ at site 2, q_{d,CO_2}	[mol/kg]	3.24
Saturation capacity of N ₂ at site 1, q_{b,N_2}	[mol/kg]	3.257
Saturation capacity of N ₂ at site 2, q_{d,N_2}	[mol/kg]	3.24
Adsorption equilibrium constant of CO ₂ at site 1, $b_{CO_2,0}$	[m ³ /mol]	2.09×10^{-7}
Adsorption equilibrium constant of CO ₂ at site 2, $d_{CO_2,0}$	[m ³ /mol]	1.06×10^{-7}
Adsorption equilibrium constant of N ₂ at site 1, $b_{N_2,0}$	[m ³ /mol]	7.96×10^{-7}
Adsorption equilibrium constant of N ₂ at site 2, $d_{N_2,0}$	[m ³ /mol]	6.94×10^{-7}
Molar internal energy change for adsorption of CO ₂ at site 1, $-\Delta U_{b,CO_2}$	[kJ/mol]	42.67
Molar internal energy change for adsorption of CO ₂ at site 2, $-\Delta U_{d,CO_2}$	[kJ/mol]	32.21
Molar internal energy change for adsorption of N ₂ at site 1, $-\Delta U_{b,N_2}$	[kJ/mol]	18.86
Molar internal energy change for adsorption of N ₂ at site 2, $-\Delta U_{d,N_2}$	[kJ/mol]	17.76

Table 3
Design variable bounds for the optimisation of the 4-step PVSA process.

Variable	t_{ads}	P_H	v_{feed}	t_{bd}	P_L	t_{vac}
Unit	s	bar	m s ⁻¹	s	bar	s
Lower Bound	20	1	0.1	30	0.07	30
Upper Bound	100	10	2	200	3	200

of the PVSA process, the low pressure set point was fixed at 0.05 bar and the pressurisation time at 20 s.

The two optimisation problems described above are solved using the IPOPT nonlinear programming algorithm (Wächter and Biegler, 2006). Although IPOPT is employed here, the overall optimisation framework is readily transferable to other gradient-based optimisers and thus the methodology is not restricted to a single numerical solver implementation. Furthermore, due to the nonconvexity of the optimisation problem, this gradient-based algorithm could converge to local minima. To explore this, a multi-start is conducted for each optimisation problem, where ten runs of the algorithm are completed, each

from a random start point within the design space. The random point is not required to satisfy the constraints. The gradient-based approach is compared to the NSGA-II algorithm, which is widely used in the adsorption literature for the multi-objective black-box optimisation, where a population size of 72 with 70 generations was used. Sobol sampling of the design space is also used for this comparison with 4096 samples. All the simulations were performed on an Intel Core i5-10500 CPU with 8 GB of RAM.

4. Results

The performance of the end-to-end differentiable model is quantified at various levels. First, we compare the simulation results with those from a similar model within the literature. Second, we compare by means of forward simulations the direct determination method to accelerate the convergence to CSS to the successive substitution method. Lastly, we compare results obtained by solving various multi-objective optimisation problems.

4.1. Simulation verification

The framework developed in this work is based upon the standard model formulation used throughout the adsorption literature (Haghpanah et al., 2013; Ward and Pini, 2022a; Wilkins et al., 2025; Pai et al., 2025; Peh et al., 2022; Subraveti et al., 2026, 2019). Before the methods introduced in this work are compared, the model is verified against published results to confirm that the implementation produces consistent predictions. The minimum energy designs reported by Haghpanah et al. (2013) are used for this comparison, where constraints of 90% purity and 90% recovery are enforced. The operating conditions for each design are taken directly from the reported optimal solutions and evaluated in the present model without re-optimisation. The model parameters used for this comparison are provided in Section S9 of the Supporting Information.

Table 4 compares the productivity and specific energy consumption obtained from the present model with the values reported by Haghpanah et al. (2013) across five minimum energy designs at different evacuation pressures. The two models show good agreement, with deviations in productivity within $\pm 5\%$ and deviations in energy consumption within 4%. The purity and recovery values obtained from the present model are consistent with the enforced constraints. Minor differences between the two implementations exist in the model formulation and in the calculation of the key performance indicators (Section S9 of the Supporting Information).

4.2. Acceleration to CSS

The convergence performance to CSS of the successive substitution and direct determination methods is compared for two different initial conditions in Fig. 5: a bed saturated with the feed gas and a bed saturated with N₂ (the light component). The successive substitution method (Fig. 5a) requires 82 and 115 iterations to reach the CSS tolerance from a bed saturated with feed and with N₂, respectively. In both scenarios, linear convergence is exhibited, as expected for this fixed point iteration, and nearly 300 iterations are required to reach the error associated with the solver tolerance. In contrast, the direct determination method (Fig. 5b) requires significantly fewer iterations to attain CSS (between 5 and 10 iterations, depending on the scenario), reflecting the near-quadratic convergence behaviour expected for a Newton iteration near the solution. The scenario with the bed saturated with the feed gas as the initial condition features the best convergence, with just 5 iterations required to reach CSS. The scenarios with the bed saturated with N₂ show a slightly lower performance, because either the damping fallback method or the successive substitution method had to be deployed as fall-back options when the (undamped) Newton iteration predicted negative values of mole fraction and adsorbed

Table 4

Comparison of the present model with the results of Haghpanah et al. (2013) for the minimum energy designs at different evacuation pressures. Productivity is given in $\text{mol m}^{-3} \text{s}^{-1}$ and energy consumption in kWh t^{-1} . Δ denotes the relative difference between the present model and the reference values.

P_L [bar]	Purity	Recovery	Productivity			Energy consumption		
			This work	Ref.	Δ [%]	This work	Ref.	Δ [%]
0.02	0.906	0.894	0.297	0.312	-4.97	143.6	149.0	-3.62
0.03	0.906	0.907	0.508	0.512	-0.87	194.5	201.5	-3.46
0.05	0.903	0.909	0.333	0.348	-4.42	289.3	298.9	-3.23
0.07	0.910	0.914	0.704	0.688	+2.35	373.2	388.7	-3.98
0.10	0.918	0.911	0.721	0.694	+3.86	486.8	499.7	-2.56




Fig. 5. Convergence plots for the successive substitution (left) and direct determination (right) methods for two initial conditions: feed- (solid line) and nitrogen-saturated bed (dashed line). In the direct determination method, two different fallback methods were used when negative mole fractions or adsorbed quantities were encountered successive substitution (dashed line) and damped Newton (dotted-dashed line). Design parameters: $P_H = 8$ bar, $P_L = 1.5$ bar, $v_{\text{feed}} = 0.8 \text{ ms}^{-1}$, $t_{\text{ads}} = 50$ s, $t_{\text{bd}} = 100$ s, $t_{\text{evac}} = 100$ s.

quantity for the next iteration. A similar observation was reported by Croft and LeVan (1994a) in their air purification example where the initial state is an empty bed. In our study, the damping fallback method provides an improved convergence to the CSS than reverting to the successive substitution method.

The two methods are again compared in Fig. 6 by considering 4096 different designs in the six-dimensional design space (Table 3). To this end, we have considered a cycle that begins with the high pressure adsorption step (the bed is saturated with the feed gas). In the figure, the violin plots show the number of iterations required to reach CSS for all 4096 samples and for the two methods (the sample set was identical for both methods). The direct determination method requires on average just 5.1 iterations to reach CSS with all of the 4096 designs converging within 11 iterations. The successive substitution method requires on average 112 cycles to reach CSS, with a much wider distribution of number of iterations (40–2000) compared to the direct determination method.

The mass balance error (Section S8.5 in the Supporting Information) of the numerical scheme was evaluated at CSS, as well as the

difference between the KPIs obtained for the two methods are given in Table 5. Both methods feature negligibly small errors in the mass balance. Across all 4096 designs, there is excellent agreement of the KPIs computed through both methods, with all designs being within 0.1% of each other for the four KPIs recorded. This shows that the two methods converge to the same CSS for all the designs explored.

It is possible that a process cycle can exhibit multiple periodic states or have a periodic state that repeats every 2 cycles, but this was not observed for this process (Croft and LeVan, 1994b). Furthermore, it is possible the direct determination method could converge to a point that is mathematically a CSS, but physically unattainable. A CSS of this type is unstable such that any disturbance would cause the system to diverge from this CSS. Mathematically, this CSS would have a spectral radius greater than 1. The spectral radius for the design corresponding to the convergence plots in Fig. 5 is 0.9509, confirming a stable and physically plausible CSS. The value of this spectral radius also gives an insight into the cycle dynamics where values closer to unity exhibit slower cycle dynamics and would represent a process that requires more cycles to reach CSS.




Fig. 6. Violin plots of the number of iterations required to reach CSS for the direct determination and successive substitution methods across 4096 Sobol sampled designs. The white box illustrates the quartiles of the distribution and the whiskers correspond to 1.5 times the interquartile range. The mean of the distribution is indicated by the diamond symbol.

Table 5

Mass balance errors and KPI differences between the successive substitution and direct determination methods.

	Successive Substitution		Direct Determination	
	Mean (%)	Max (%)	Mean (%)	Max (%)
<i>Mass Balance Error</i>				
CO ₂	6.97e-03	0.830	1.73e-03	0.148
N ₂	1.81e-03	4.89e-02	3.33e-04	5.21e-03
<i>KPI Difference vs. Successive Substitution^a</i>				
Purity	—	—	9.31e-04	2.51e-02
Recovery	—	—	7.83e-04	6.99e-03
Productivity	—	—	1.36e-03	5.59e-02
Energy consumption	—	—	7.94e-04	3.96e-03

^a KPI differences are computed as the absolute percentage deviation of the Direct Determination method relative to Successive Substitution as the reference solution.

4.3. PVSA process optimisation

4.3.1. Three decision variables

The direct determination method provides acceleration to the solution speed of each individual simulation for use within an optimisation framework. For this design problem, three different approaches are compared: Sobol sampling of the design space (4096 samples), the NSGA-II algorithm and the gradient-based optimisation using the IPOPT algorithm. The latter exploits the gradient information readily available from the differentiable model. The comparison is first performed for a reduced design problem involving three decision variables: the gas feed velocity, v_{feed} , the high pressure set point, P_H , and the intermediate pressure set point, P_I . The adsorption, blowdown and evacuation step durations are fixed at 50, 100 and 100 s, respectively (Sachio et al., 2024).

Fig. 7 shows the Pareto fronts obtained for the multi-objective optimisation problem where both purity and recovery are maximised (Eq. (17)). To construct the gradient-based Pareto front, the epsilon constraint on purity is varied between 90% and 99.5% in 0.5% increments, with additional evaluations at 99.9% and 99.99% purity. The three solution methods identify similar Pareto fronts and capture the




Fig. 7. Pareto front for the purity-recovery optimisation problem solved with Sobol sampling, the NSGA-II algorithm and the IPOPT algorithm which uses gradients from the differentiable model. For the latter, a multi-start strategy (10 random initial start points) is adopted and all runs are shown with the green crosses.

trade-off between CO₂ purity and CO₂ recovery. The gradient-based optimisation yields a slightly improved front than the genetic algorithm, with both methods improving upon Sobol sampling, which produces a cloud of feasible designs lying close to but below the optimal front. The multi-start strategy adopted for the gradient-based optimisation reveals that the non-convexity of this problem is significant in the sub-problem with the 98% purity constraint. Two of the ten multi-start points converge to a distinct sub-optimal point, characterised by a larger intermediate pressure and feed velocity, while the high pressure converges to its upper bound in all cases. The translucent green crosses in Fig. 7 indicate these locally optimal points. The raw data for all gradient-based optimisation runs are provided in the Supplementary Information.

Fig. 8 shows the Pareto fronts obtained for the multi-objective optimisation problem where productivity is maximised and energy usage is minimised (Eq. (18)) - again, for the three methods. The genetic algorithm and the gradient-based optimisation produce Pareto fronts that are in close agreement, while the approach that uses the 4096 Sobol samples demonstrates obvious deficiencies. In fact, as shown in Fig. 7, relatively few points in the whole design space satisfy simultaneously the constraints on CO₂ purity (95%) and CO₂ recovery (90%). In other words, for the same number of Sobol samples, a reduction in the size of the sampled input space would be needed to increase the frequency of feasible designs. In contrast to the purity/recovery optimisation problem, every multi-start point for each epsilon-constrained sub-problem converges to the same solution. The added purity and recovery constraints restrict the feasible region for this optimisation which likely reduces the possibility of converging to multiple optima. The gradient-based algorithm is also used to evaluate single-objective optimisation problems for maximising the productivity (indicated by the red cross) and for minimising the energy consumption (indicated by the purple cross) and to explore the extreme ends of the Pareto front. We note that neither the genetic algorithm, nor the Sobol sampling points reached the point with the lowest energy consumption, highlighting the value of the single-objective gradient-based optimisation to capture the extremes of the Pareto curve. The convergence of the multi-start to identical points in the gradient-based optimisation and the




Fig. 8. Pareto front for the minimisation of energy consumption and maximisation of productivity with 95% purity and 90% recovery constraints comparing Sobol sampling, the NSGA-II algorithm and the IPOPT algorithm which uses gradients from the differentiable model. Scenarios with three decision variables.




Fig. 9. Pareto front for the minimisation of energy consumption and maximisation of productivity with 95% purity and 90% recovery constraints comparing Sobol sampling, the NSGA-II algorithm and the IPOPT algorithm which uses gradients from the differentiable model. Scenarios with six decision variables.

strong agreement with the NSGA-II optimiser provide confidence that the identified Pareto front closely represents the true optimal trade-off for this three-variable problem.

4.3.2. Six decision variables

We explore in this section the full design problem to include all six decision variables in Table 3. Fig. 9 shows the Pareto fronts obtained with the same three methods as in the previous section. For this higher-dimensional problem, the gradient-based optimisation produces

a Pareto front that is significantly improved over the one produced by the genetic algorithm. Notably, the latter also shows a shorter Pareto curve and does not extend into the region of low and high productivity bounded by the solution of the two single-objective optimisation problems. Again – and not surprisingly – both the gradient-based algorithm and the genetic algorithm outperform Sobol sampling.

The widening of the gap between these methods relative to the case with three decision variables illustrates the curse of dimensionality in this optimisation problem. The design space grows exponentially with the number of decision variables. Both the sampling- and population-based methods must explore this space with finite number of evaluations. With 4096 Sobol samples in a six-dimensional space, the coverage is sparse and is equivalent to approximately four levels per variable if the sampling were uniform ($4^6 = 4096$). The gradient-based approach is less affected by this scaling as the gradient information available to the optimiser grows linearly with the number of decision variables, which provides directional information that guides the search. A further distinction between the three methods is in the termination criterion. The gradient-based optimisation provides a first-order optimality guarantee at convergence, terminating when the gradients of the objective and constraints satisfy the necessary conditions for a local optimum. The genetic algorithm and Sobol sampling, by contrast, are typically run with a predefined computational budget and offer no such guarantee. The limitation of the gradient-based approach is the potential for convergence to local optima, which can be mitigated through a multi-start strategy, though global optimality cannot be guaranteed.

To further compare the solutions obtained with the gradient-based optimisation and the genetic algorithm, Fig. 10 shows the values of the decision variables (normalised to their bounds) of the optimal designs as a function of the productivity (i.e. moving along the Pareto front). The solid lines refer to the solution of each epsilon-constrained sub-problem and the translucent crosses represent results from all 10 multi-start solutions for each sub-problem. The results from the NSGA-II algorithm are shown by the scattered translucent circles. Overall, there is general agreement between the trends outlined by the two optimisation algorithms. The shift of the results obtained with the genetic algorithm and their limited coverage of productivity values reflect the deficiencies observed in Fig. 9 when this algorithm is compared to the solution obtained by gradient-based optimisation. As discussed above, these differences reflect the lower precision of the population-based search within its computational budget.

The results from the multi-start for the gradient-based optimisation provide further insight into the optimisation landscape. For the majority of the Pareto front, all 10 starts converge to the same solution in both the objective and decision variables space. At the low energy and low productivity end of the Pareto front, the optimisation landscape is very flat, as can be seen from the large range of productivities that yield a very similar energy consumption in Fig. 9. In this region, both the high pressure and the intermediate pressure converge to the same points. These variables contribute significantly to the energy consumption. The adsorption time and the feed velocity, however, freely trade off against each other, as their product determines the total feed volume per adsorption step. This observation suggests that the contribution of the pressure drop to the overall energy consumption is smaller than the optimality tolerance, and the optimiser does not resolve a unique combination of feed velocity and adsorption time. At the high productivity end of the Pareto front, local optima emerge. At the third highest productivity point on the front ($14.6 \text{ tonneCO}_2 \text{ m}^{-3} \text{ day}^{-1}$), 7 of the 10 starts converge to the best solution, while 3 converge to a local optimum that differs by only 0.2 kWh/tonne in energy consumption. Yet, these solution differ significantly in their decision variables, suggesting that there are different arcs of optimal solutions within this region of the Pareto front. At the maximum productivity extreme, 2 of the 10 starts find a local optimum that is dominated by the other 8 solutions. In both scenarios where local optima are found, the impact on the objective function is small.




Fig. 10. Decision variables of the optimal designs obtained by gradient-based optimisation (solids lines) and by the NSGA-II algorithm (symbols), i.e. the Pareto curves in Fig. 9. The decision variables are normalised to their bounds and are plotted as a function of the productivity obtained for those points.

5. Discussion

The direct determination method requires significantly fewer iterations to attain CSS than successive substitution. However, each iteration of the direct determination method incurs additional computation for the evaluation of the cycle Jacobian ($\partial x_k / \partial x_k^0$ in Eq. (4)), such that a comparison of the total computational time is necessary. Furthermore, a comparison is made to the successive substitution method for a near-identical model implemented in MATLAB, which represents the standard practice for this model formulation. The two implementations differ in the solution method: the JAX formulation uses automatic differentiation for the Jacobian evaluation of the ODE right-hand side function within the implicit solver, rather than numerical approximation in the MATLAB version, and benefits from just-in-time compilation. As shown in Fig. 11 for 4096 Sobol sampled designs, the direct determination method in JAX requires an average wall-clock time of 4.54 s to reach CSS, compared to 14.78 s for successive substitution in JAX and approximately 20 times longer for the MATLAB implementation with successive substitution. These times do not include the compilation time for the JAX implementation, which is 26 s for the direct determination method. This is a one-time cost, as once compiled, all subsequent designs are evaluated without recompilation.

The gradient-based optimisation method outperforms both Sobol sampling and the NSGA-II algorithm in attaining the Pareto front for the multi-objective optimisation problems studied. Table 6 compares the number of model evaluations and the total computational time for each method applied to the productivity/energy usage design problem shown in Figs. 8 and 9. Again, we also compare in Table 6 solutions obtained using the differentiable model in JAX and the conventional model in MATLAB. The solution method in terms of the finite volume scheme, the ODE solver and the optimisation structure is the same for the model in JAX and the model in MATLAB, with the former providing computation of gradients to improve the computational speed and the Pareto front obtained. We acknowledge that the direct comparison of the computational time remain difficult as the methods differ both in their structure and in the quality of the solution. Yet, these values




Fig. 11. Violin plots of the computational time required to reach CSS for the direct determination and successive substitution methods in JAX as well as the successive substitution method in MATLAB across 4096 Sobol sampled designs. The white box illustrates the quartiles of the distribution and the whiskers correspond to 1.5 times the interquartile range. The mean of the distribution is indicated by the diamond symbol.

provide useful context for the practitioner. The Sobol sampling method evaluates 4096 points to attain both the maximal purity and recovery front and the productivity/energy Pareto front, unlike a targeted optimisation problem. The GA, with a population size of 72 and with 70 generations, evaluates the model 5040 times for a single optimisation problem.

In the gradient-based optimisation, the problem is broken into single-objective sub-problems as well as having a multi-start to evaluate the non-convexity of the problem, such that the comparison is not like for like. For the purity/recovery optimisation problem, a total of 130 sub-problems are completed to obtain the Pareto front with the

Table 6

The performance of the differentiable model in JAX is evaluated across different optimisation methods and against the conventional benchmark model.

	Differentiable Model (JAX)				Conventional Model (MATLAB)
	IPOPT (3 var.)	IPOPT (6 var.)	NSGA-II	Sobol Sampling	NSGA-II
Model evaluations ^a	167	315	5040	4096	5040
Wall-clock time	9.5 min	16 min	7.1 hr	5.2 hr	129 hr
Speedup		×486	×18		×1

^a IPOPT evaluation counts are for 12 Pareto front points without multistart

multistart strategy. Each individual sub problem requires an average of 18 model evaluations and 188 s of wall-clock time to complete. The statistics for the productivity/energy consumption optimisation problems for both the three variable optimisation and the six variable optimisation are given in Table 6. For these problems, a total of 120 sub-problems are completed. An average of 13 model evaluations and a wall-clock time of 47 s are required to solve each sub-problem in the three dimensional optimisation, and 23 evaluations and 80 s for each sub-problem in the six variable optimisation. A multi-start strategy is not deemed essential for the productivity/energy consumption optimisation problem as the nonconvexity of the design space does not affect the obtained Pareto, thus it is not included in the statistics in Table 6. The statistics therefore just include 12 sub-problems for evaluating the Pareto front. As such, the gradient-based optimisation with the differentiable framework yields a speed-up of over 2 orders of magnitude compared to the state-of-the-art representative approach.

We note that parallelism offers a direct route to reducing the computational time of adsorption process optimisation, as many of the computational workflows discussed involve independent model evaluations that can be distributed across multiple processors. The extent to which each method benefits from parallelisation, however, differs. Sobol sampling is trivially parallel, as all 4096 designs are independent and can be evaluated simultaneously given sufficient computational resources. The NSGA-II genetic algorithm is parallel within each generation, where the population of candidate designs can be evaluated concurrently, though synchronisation is required between generations. Each individual optimisation run for the gradient-based optimisation is inherently sequential, however the parallelism potential comes with the potential to run each sub-problem in the multi-objective optimisation in parallel. Similarly, if a multi-start strategy is desired, these can be run in parallel.

6. Conclusion

The design of cyclic adsorption processes presents a significant computational challenge. The requirement to simulate from startup to cyclic steady state for every candidate design, combined with the use of derivative-free optimisation, results in design campaigns requiring hundreds to thousands of CPU hours. This burden limits the ability to systematically explore the expanding design space of new adsorbent materials, cycle configurations, and operating conditions. In this work, we have demonstrated that differentiable programming addresses this challenge by transforming the entire process model into an end-to-end differentiable computational workflow, while retaining the same physical model and numerical discretisation used throughout the adsorption literature. A four-step pressure vacuum swing adsorption process for post-combustion carbon capture was implemented in JAX, with automatic differentiation providing exact gradients throughout the model without the manual derivation of sensitivity equations.

The framework delivers acceleration at three levels of the design workflow. At the level of the ODE solution, just-in-time compilation and automatic differentiation of the implicit solver reduce the simulation time to cyclic steady state from 91 s of wall-clock time in MATLAB to 15 s in JAX for the successive substitution method. At the level of

cyclic steady state convergence, the differentiation of a single process cycle provides the Jacobian required for a Newton-based direct determination method. This exploits the near-quadratic convergence of the Newton iteration, which requires an average of 5.1 iterations to reach cyclic steady state compared to 112 for successive substitution, further reducing the simulation time to 4.5 s — a reduction of approximately 20-fold relative to the conventional MATLAB implementation. At the level of the optimisation, the availability of exact gradients of the key performance indicators with respect to the design variables enables the use of the IPOPT algorithm for multi-objective optimisation via the epsilon-constraint method. For the six-variable problem maximising productivity and minimising energy consumption subject to purity and recovery constraints, the gradient-based approach obtains the Pareto front in approximately 16 min — over two orders of magnitude faster than the conventional genetic algorithm approach. The resulting front is also of superior quality, capturing the extremes of the trade-off curve that the genetic algorithm does not reach.

The central advantage of the approach lies in its generality. Because automatic differentiation operates throughout the computational workflow irrespective of the specific model formulation, exact gradients are made available wherever they are useful, without the manual derivation of sensitivity equations that has historically limited the adoption of gradient-based methods in the adsorption community. The framework is therefore readily applicable to alternative adsorbents, cycle configurations, and column designs. Future work will extend the approach to more complex multi-column cycles, integrate the differentiable model with GPU-accelerated solvers for the massively parallel evaluation of designs, and exploit the available gradients for uncertainty quantification and adsorbent screening. More broadly, this work illustrates how differentiable programming can shift cyclic process design from slow black-box simulation with derivative-free optimisation towards efficient, gradient-enhanced modelling and optimisation, with relevance to the wider class of periodic processes encountered in chemical engineering.

CRedit authorship contribution statement

Alex Glover: Writing – review & editing, Writing – original draft, Visualization, Software, Methodology, Formal analysis, Conceptualization. **Maria M. Papathanasiou:** Writing – review & editing, Supervision, Conceptualization. **Ronny Pini:** Writing – review & editing, Supervision, Resources, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

A.G. gratefully acknowledges funding from the EPSRC DTP scholarship granted to the Department of Chemical Engineering at Imperial College London and by the Bansal Bursary. The authors also thank Ben Moseley and Ashwin Kumar Rajagopalan, along with members of their respective research groups, for valuable discussions.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.compchemeng.2026.109904>.

Data availability

Data will be made available on request.

References

- Agarwal, A., Biegler, L.T., Zitney, S.E., 2010. A superstructure-based optimal synthesis of PSA cycles for post-combustion CO₂ capture. *AIChE J.* 56 (7), 1813–1828. <http://dx.doi.org/10.1002/aic.12107>.
- Alsubeih, M., Jessop, A., Moseley, B., Fonte, C.P., Rajagopalan, A.K., 2025. Modern, efficient, and differentiable transport equation models using JAX: Applications to population balance equations. *Ind. Eng. Chem. Res.* 64 (8), 4541–4553. <http://dx.doi.org/10.1021/acs.iecr.4c04208>.
- Antunes Morgado, N., Mazzotti, M., 2026. 3D modeling of isothermal adsorption in structured contactors. *Comput. Chem. Eng.* 205, 109404. <http://dx.doi.org/10.1016/j.compchemeng.2025.109404>.
- Beck, J.H., 2014. Efficient targeted optimisation for the design of pressure swing adsorption systems for CO₂ capture in power plants.
- Bezgin, D.A., Buhendwa, A.B., Adams, N.A., 2023. JAX-fluids: A fully-differentiable high-order computational fluid dynamics solver for compressible two-phase flows. *Comput. Phys. Comm.* 282, 108527. <http://dx.doi.org/10.1016/j.cpc.2022.108527>.
- Biegler, L.T., Jiang, L., Fox, V.G., 2005. Recent advances in simulation and optimal design of pressure swing adsorption systems. *Sep. Purif. Rev.* 33 (1), 1–39. <http://dx.doi.org/10.1081/SPM-120039562>.
- Chatziasteriou, C.C., Georgiadis, M.C., Kikkinides, E.S., 2025. Surrogate modeling and optimization of pressure/vacuum swing adsorption (P/VSA) processes for carbon capture from post-combustion CO₂ point sources. *Chem. Eng. Res. Des.* 213, 286–295. <http://dx.doi.org/10.1016/j.cherd.2024.12.002>.
- Croft, D.T., LeVan, M., 1994a. Periodic states of adsorption cycles—I. direct determination and stability. *Chem. Eng. Sci.* 49 (11), 1821–1829. [http://dx.doi.org/10.1016/0009-2509\(94\)80067-7](http://dx.doi.org/10.1016/0009-2509(94)80067-7).
- Croft, D.T., LeVan, M., 1994b. Periodic states of adsorption cycles—II. solution spaces and multiplicity. *Chem. Eng. Sci.* 49 (11), 1831–1841. [http://dx.doi.org/10.1016/0009-2509\(94\)80068-5](http://dx.doi.org/10.1016/0009-2509(94)80068-5).
- Cruz, P., Magalhães, F.D., Mendes, A., 2005. On the optimization of cyclic adsorption separation processes. *AIChE J.* 51 (5), 1377–1395. <http://dx.doi.org/10.1002/aic.10400>.
- Deb, K., Agrawal, S., Pratap, A., Meyarivan, T., 2000. A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: NSGA-II. In: Schoenauer, M., Deb, K., Rudolph, G., Yao, X., Lutton, E., Merelo, J.J., Schwefel, H.-P. (Eds.), *Parallel Problem Solving from Nature PPSN VI*. Springer, Berlin, Heidelberg, pp. 849–858. http://dx.doi.org/10.1007/3-540-45356-3_83.
- Ding, Y., LeVan, M.D., 2001. Periodic states of adsorption cycles III. convergence acceleration for direct determination. *Chem. Eng. Sci.* 56 (17), 5217–5230. [http://dx.doi.org/10.1016/S0009-2509\(01\)00127-0](http://dx.doi.org/10.1016/S0009-2509(01)00127-0).
- Douglas LeVan, M., Croft, D.T., 1995. Determination of periodic states of pressure swing adsorption cycles. *Gas Sep. Purif.* 9 (1), 13–16. [http://dx.doi.org/10.1016/0950-4214\(95\)92171-8](http://dx.doi.org/10.1016/0950-4214(95)92171-8).
- Dowling, A.W., Vetukuri, S.R.R., Biegler, L.T., 2012. Large-scale optimization strategies for pressure swing adsorption cycle synthesis. *AIChE J.* 58 (12), 3777–3791. <http://dx.doi.org/10.1002/aic.13928>.
- Estupiñan Perez, L., Sarkar, P., Rajendran, A., 2019. Experimental validation of multi-objective optimization techniques for design of vacuum swing adsorption processes. *Sep. Purif. Technol.* 224, 553–563. <http://dx.doi.org/10.1016/j.seppur.2019.05.039>.
- GDSFactory, 2026. Gdsfactory/klujax. URL <https://github.com/gdsfactory/klujax>. original-date: 2021-06-01T19:26:37Z.
- Haghighanah, R., Majumder, A., Nilam, R., Rajendran, A., Farooq, S., Karimi, I.A., Amanullah, M., 2013. Multiobjective optimization of a four-step adsorption process for postcombustion CO₂ capture via finite volume simulation. *Ind. Eng. Chem. Res.* 52 (11), 4249–4265. <http://dx.doi.org/10.1021/ie302658y>.
- Hao, Z., Caspari, A., Schweidtmann, A.M., Vaupel, Y., Lapkin, A.A., Mhamdi, A., 2021. Efficient hybrid multiobjective optimization of pressure swing adsorption. *Chem. Eng. J.* 423, 130248. <http://dx.doi.org/10.1016/j.cej.2021.130248>.
- Henrotin, A., Heymans, N., Duprez, M.-E., De Weireld, G., 2026. Modelling and optimization of vacuum pressure swing adsorption CO₂ capture pilot using MIL-160(Al). *Chem. Eng. Res. Des.* 226, 14–30. <http://dx.doi.org/10.1016/j.cherd.2025.12.030>.
- Jiang, L., Biegler, L.T., Fox, V.G., 2003. Simulation and optimization of pressure-swing adsorption systems for air separation. *AIChE J.* 49 (5), 1140–1157. <http://dx.doi.org/10.1002/aic.690490508>.
- Kawajiri, Y., Biegler, L.T., 2006. Optimization strategies for simulated moving bed and PowerFeed processes. *AIChE J.* 52 (4), 1343–1350. <http://dx.doi.org/10.1002/aic.10736>.
- Khajuria, H., Pistikopoulos, E.N., 2013. Optimization and control of pressure swing adsorption processes under uncertainty. *AIChE J.* 59 (1), 120–131. <http://dx.doi.org/10.1002/aic.13783>.
- Khurana, M., Farooq, S., 2016. Adsorbent screening for postcombustion CO₂ capture: A method relating equilibrium isotherm characteristics to an optimum vacuum swing adsorption process performance. *Ind. Eng. Chem. Res.* 55 (8), 2447–2460. <http://dx.doi.org/10.1021/acs.iecr.5b04531>.
- Khurana, M., Farooq, S., 2017. Integrated adsorbent-process optimization for carbon capture and concentration using vacuum swing adsorption cycles. *AIChE J.* 63 (7), 2987–2995. <http://dx.doi.org/10.1002/aic.15602>.
- Ko, D., Siriwardane, R., Biegler, L.T., 2003. Optimization of a pressure-swing adsorption process using zeolite 13X for CO₂ sequestration. *Ind. Eng. Chem. Res.* 42 (2), 339–348. <http://dx.doi.org/10.1021/ie0204540>.
- Krishnamurthy, S., Rao, V.R., Guntuka, S., Sharratt, P., Haghighanah, R., Rajendran, A., Amanullah, M., Karimi, I.A., Farooq, S., 2014. CO₂ capture from dry flue gas by vacuum swing adsorption: A pilot plant study. *AIChE J.* 60 (5), 1830–1842. <http://dx.doi.org/10.1002/aic.14435>.
- Latifi, M.A., Salhi, D., Tondeur, D., 2008. Optimisation-based simulation of a pressure swing adsorption process. *Adsorption* 14 (4), 567–573. <http://dx.doi.org/10.1007/s10450-008-9130-0>.
- Lee, I., Buse, S., Winfree, E., 2025. Differentiable programming of indexed chemical reaction networks and reaction-diffusion systems. *LIPICs* 347, 4:1–4:23. <http://dx.doi.org/10.4230/LIPICs.DNA.31.4>.
- LeVeque, R.J., 2002. *Finite Volume Methods for Hyperbolic Problems*. Cambridge University Press, Google-Books-ID. QazcnD7GUoUC.
- Liao, Y., Wright, A., Li, J., 2025. Simulation and optimisation of vacuum (pressure) swing adsorption with simultaneous consideration of real vacuum pump data and bed fluidisation. *Sep. Purif. Technol.* 358, 130354. <http://dx.doi.org/10.1016/j.seppur.2024.130354>.
- Mahajan, S., Lahtinen, M., 2022. Recent progress in metal-organic frameworks (MOFs) for CO₂ capture at different pressures. *J. Environ. Chem. Eng.* 10 (6), 108930. <http://dx.doi.org/10.1016/j.jece.2022.108930>.
- Marcinek, A., Guderian, J., Bathen, D., 2020. Performance determination of high-purity N₂-PSA-plants. *Adsorption* 26 (7), 1215–1226. <http://dx.doi.org/10.1007/s10450-020-00204-9>.
- Marx, D., Joss, L., Hefti, M., Mazzotti, M., 2016. Temperature swing adsorption for postcombustion CO₂ capture: Single- and multicolumn experiments and simulations. *Ind. Eng. Chem. Res.* 55 (5), 1401–1412. <http://dx.doi.org/10.1021/acs.iecr.5b03727>.
- Miettinen, K., 1998. Nonlinear multiobjective optimization. In: *International Series in Operations Research & Management Science*, vol. 12, Springer US, Boston, MA, <http://dx.doi.org/10.1007/978-1-4615-5563-6>.
- Nilchan, S., Pantelides, C., 1998. On the optimisation of periodic adsorption processes. *Adsorption* 4 (2), 113–147. <http://dx.doi.org/10.1023/A:1008823102106>.
- Pai, K.N., Haghighanah, R., Edsall, W., 2025. Superstructure model for the simultaneous design and optimization of a pressure vacuum swing adsorption process cycle. *Ind. Eng. Chem. Res.* 64 (1), 627–638. <http://dx.doi.org/10.1021/acs.iecr.4c02842>.
- Pai, K.N., Prasad, V., Rajendran, A., 2020a. Experimentally validated machine learning frameworks for accelerated prediction of cyclic steady state and optimization of pressure swing adsorption processes. *Sep. Purif. Technol.* 241, 116651. <http://dx.doi.org/10.1016/j.seppur.2020.116651>.
- Pai, K.N., Prasad, V., Rajendran, A., 2020b. Generalized, adsorbent-agnostic, artificial neural network framework for rapid simulation, optimization, and adsorbent screening of adsorption processes. *Ind. Eng. Chem. Res.* 59 (38), 16730–16740. <http://dx.doi.org/10.1021/acs.iecr.0c02339>.
- Peh, S.B., Farooq, S., Zhao, D., 2022. A metal-organic framework (MOF)-based temperature swing adsorption cycle for postcombustion CO₂ capture from wet flue gas. *Chem. Eng. Sci.* 250, 117399. <http://dx.doi.org/10.1016/j.ces.2021.117399>.
- Pini, R., Petit, C., Danaci, D., Haghighanah, R., Luberti, M., Ribeiro, A.M., Subraveti, S.G., Ward, A., 2026. CO₂ capture by adsorption: Research progress and technology demonstration. *Int. J. Greenh. Gas Control.* 153, 104652. <http://dx.doi.org/10.1016/j.jggc.2026.104652>.
- Raganati, F., Miccio, F., Ammendola, P., 2021. Adsorption of carbon dioxide for post-combustion capture: A review. *Energy Fuels* 35 (16), 12845–12868. <http://dx.doi.org/10.1021/acs.energyfuels.1c01618>.
- Rajagopalan, A.K., Avila, A.M., Rajendran, A., 2016. Do adsorbent screening metrics predict process performance? A process optimisation based study for post-combustion capture of CO₂. *Int. J. Greenh. Gas Control.* 46, 76–85. <http://dx.doi.org/10.1016/j.jggc.2015.12.033>.
- Ruthven, D.M., 1984. *Principles of Adsorption and Adsorption Processes*. John Wiley & Sons, Google-Books-ID. u7wq21njR3UC.
- Sachio, S., Ward, A., Pini, R., Papatthanasou, M.M., 2024. Operability-economics trade-offs in adsorption-based CO₂ capture processes. *Commun. Eng.* 3 (1), 1–10. <http://dx.doi.org/10.1038/s44172-024-00244-x>.
- Sapienza, F., Bolibar, J., Schäfer, F., Groenke, B., Pal, A., Boussange, V., Heimbach, P., Hooker, G., Pérez, F., Persson, P.-O., Rackauckas, C., 2024. Differentiable programming for differential equations: A review. <http://dx.doi.org/10.48550/arXiv.2406.09699>.

- Schoenholz, S.S., Cubuk, E.D., 2021. JAX, M.D. A framework for differentiable physics*. *J. Stat. Mech. Theory Exp.* 2021 (12), 124016. <http://dx.doi.org/10.1088/1742-5468/ac3ae9>.
- Shampine, L.F., Reichelt, M.W., 1997. The MATLAB ODE suite. *SIAM J. Sci. Comput.* 18 (1), 1–22. <http://dx.doi.org/10.1137/S1064827594276424>.
- Sircar, S., 1992. Estimation of isosteric heats of adsorption of single gas and multicomponent gas mixtures. *Ind. Eng. Chem. Res.* 31 (7), 1813–1819. <http://dx.doi.org/10.1021/ie00007a030>.
- Subramanian Balashankar, V., Rajagopalan, A.K., De Pauw, R., Avila, A.M., Rajendran, A., 2019. Analysis of a batch adsorber analogue for rapid screening of adsorbents for postcombustion CO₂ capture. *Ind. Eng. Chem. Res.* 58 (8), 3314–3328. <http://dx.doi.org/10.1021/acs.iecr.8b05420>.
- Subraveti, S.G., Karimi, K., Gazzani, M., 2026. Vector acceleration methods for faster convergence of cyclic steady state in adsorption process simulations. *Ind. Eng. Chem. Res.* <http://dx.doi.org/10.1021/acs.iecr.6c00407>.
- Subraveti, S.G., Karimi, K., Gazzani, M., Anantharaman, R., 2024. A Study on Accelerated Convergence of Cyclic Steady State in Adsorption Process Simulations. Breckenridge, Colorado, USA, pp. 215–221. <http://dx.doi.org/10.69997/sect.137508>.
- Subraveti, S.G., Li, Z., Prasad, V., Rajendran, A., 2019. Machine learning-based multiobjective optimization of pressure swing adsorption. *Ind. Eng. Chem. Res.* 58 (44), 20412–20422. <http://dx.doi.org/10.1021/acs.iecr.9b04173>.
- Ternet, D.J., Biegler, L.T., 1998. Recent improvements to a multiplier-free reduced hessian successive quadratic programming algorithm. *Comput. Chem. Eng.* 22 (7–8), 963–978. [http://dx.doi.org/10.1016/S0098-1354\(98\)00009-X](http://dx.doi.org/10.1016/S0098-1354(98)00009-X).
- Todd, R.S., He, J., Webley, P.A., Beh, C., Wilson, S., Lloyd, M.A., 2001. Fast finite-volume method for PSA/VSA cycle Simulation Experimental validation. *Ind. Eng. Chem. Res.* 40 (14), 3217–3224. <http://dx.doi.org/10.1021/ie0008070>.
- Tsay, C., Pattison, R.C., Baldea, M., 2018. A pseudo-transient optimization framework for periodic processes: Pressure swing adsorption and simulated moving bed chromatography. *AIChE J.* 64 (8), 2982–2996. <http://dx.doi.org/10.1002/aic.15987>.
- Vetukuri, S.R.R., Biegler, L.T., Walther, A., 2010. An inexact trust-region algorithm for the optimization of periodic adsorption processes. *Ind. Eng. Chem. Res.* 49 (23), 12004–12013. <http://dx.doi.org/10.1021/ie100706c>.
- Wächter, A., Biegler, L.T., 2006. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Math. Program.* 106 (1), 25–57. <http://dx.doi.org/10.1007/s10107-004-0559-y>.
- Ward, A., Pini, R., 2022a. Efficient Bayesian optimization of industrial-scale pressure-vacuum swing adsorption processes for CO₂ capture. *Ind. Eng. Chem. Res.* 61 (36), 13650–13668. <http://dx.doi.org/10.1021/acs.iecr.2c02313>.
- Ward, A., Pini, R., 2022b. Integrated uncertainty quantification and sensitivity analysis of single-component dynamic column breakthrough experiments. *Adsorption* 28 (3), 161–183. <http://dx.doi.org/10.1007/s10450-022-00361-z>.
- Ward, A., Pini, R., 2024. Design and performance evaluation of multisorbent vacuum-swing adsorption processes for postcombustion carbon capture. *Ind. Eng. Chem. Res.* 63 (31), 13787–13800. <http://dx.doi.org/10.1021/acs.iecr.4c01345>.
- Wilkins, N.S., McKinnon, D.M., Baliga, C., Nguyen, T.T.T., Moroy, B., Xu, I., Lin, X., Tulpar, O., Wheeler, N., Tran, V., Moghaddam, A.L., Cronin, D., Rezaei, S., Liu, A., 2025. Pilot-scale demonstration of a laminate-based CALF-20 rapid cycle pressure-vacuum swing adsorption process for carbon capture. <http://dx.doi.org/10.26434/chemrxiv-2025-zz73j>.
- Wilkins, N.S., Rajendran, A., 2019. Measurement of competitive CO₂ and N₂ adsorption on zeolite 13X for post-combustion CO₂ capture. *Adsorption* 25 (2), 115–133. <http://dx.doi.org/10.1007/s10450-018-00004-2>.