hydrogen

MDPI

Article

Physics-Informed Composition-to-Tank Modelling of Metal Hydride Hydrogen Storage

Amandyk Kartbayev [1,*], Dababrata Chowdhury [2] and Fuad Nasirli [3]

[1] Computer Technologies Department, Computer Science and AI Faculty, Yessenov Caspian University of Technology and Engineering, Aktau 130000, Kazakhstan
[2] Management and Entrepreneurship Department, Ravensbourne University, London SE10 0EW, UK; d.chowdhury@rave.ac.uk
[3] Computer Engineering Department, Business Administration Faculty, Azerbaijan State Oil and Industry University, Baku AZ1010, Azerbaijan; fuad.nasirli@asoiu.edu.az
* Correspondence: a.kartbayev@gmail.com

Abstract

Composition-based machine-learning models can estimate metal hydride thermodynamics, but their errors are reported at the material level, whereas storage systems are judged at the tank level, and this study connects the two. Gradient-boosting models trained on 420 hydrides with a van 't Hoff-consistent parameterization predict the room-temperature plateau, ln P_{eq}, with cross-validated mean absolute errors of 0.90 for AB_5 and 1.75 for AB_2 alloys. A reduced-order heat- and mass-transfer model of an open 7.8 kg AB_2 tank was calibrated on the pressure, temperature and flow time series of one test and reproduced two further tests within 3.5–5.5 °C and 0.4–1.0 bar. The tank data place the alloy's plateau at ln P_{eq} = 2.59–2.68, 2.0 units above the prediction, while the predicted enthalpy is within about 2 kJ mol^{-1}. With predicted thermodynamics, the simulated tank absorbs 60 g in 1.1 h instead of 8.8 h, peaks at 106 °C instead of 58 °C and delivers 31% less hydrogen; delivered capacity tolerates plateau errors of about −1.5/+0.25, charging behaviour only ±0.15. Composition models nevertheless rank substitution series correctly, as shown for $La_{1-x}Ce_xNi_5$ and $LaNi_{5-x}Al_x$, and combined with one measured isotherm they offer a fast route from alloy composition to tank design.

Keywords: metal hydride; hydrogen storage; energy storage; machine learning; plateau pressure; van 't Hoff; process modelling; model calibration; uncertainty propagation; Sobol sensitivity analysis; AB_2 and AB_5 alloys

1. Introduction

Storing hydrogen in reversible intermetallic hydrides remains attractive wherever volumetric density, low operating pressure and intrinsic safety matter more than gravimetric capacity, for example in stationary buffers for renewable electricity, backup power and small fuel-cell systems [1–4]. In these applications the alloy is not selected for its capacity alone. Its equilibrium (plateau) pressure must sit inside a window fixed by the rest of the system: high enough at the lowest discharge temperature to feed a fuel cell, and low enough at the highest charging temperature to be filled by an electrolyser or compressor [5–7]. Whether a given tank actually delivers its nominal capacity then depends on how fast the reaction heat can be removed or supplied, which is why heat management dominates engineering studies of metal hydride (MH) vessels [8–12].

For AB_5 and AB_2 intermetallics, the plateau pressure is tuned by partial substitution of the A-site element (e.g., La in $LaNi_5$) or of the B-site element (e.g., Ni) with other metals.

Check for updates

Academic Editor: Miroslav Variny

Received: 26 October 2026
Revised: 9 November 2026
Accepted: 17 November 2026
Published: 24 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

In $LaNi_5$, replacing La with Ce raises the plateau and replacing Ni with Al or Mn lowers it, at some cost in capacity and hysteresis [13–15]. Systems built around this logic include $La_{1-x}Ce_xNi_5$ accumulators coupled to photovoltaic electrolysis and fuel cells [16], and $LaNi_{5-x}Al_x$ units that absorb near-atmospheric hydrogen from solid-oxide electrolysers and release it at a few bar on heating [17,18]. In each case, the composition was chosen by measuring a short series of candidate alloys. The question this paper asks is whether part of that experimental screening can be done in silico, and with what confidence.

Data-driven models of hydride thermodynamics have made this question concrete. The HydPARK compilation and its machine-learning-ready derivative ML-HydPARK provide formation enthalpies, entropies and plateau pressures for several hundred intermetallics [19]. Models trained on composition-derived descriptors predict these quantities with useful accuracy [20–22], have been used to screen high-entropy hydrides [23], and continue to improve as databases grow [24,25]. These studies, however, report accuracy at the material level: a mean absolute error in ΔH or in $\ln P_{eq}$. A tank designer needs a different number: how wrong the predicted delivered capacity, charging time or bed temperature will be when the predicted thermodynamics are used in place of a measured isotherm. An error of one unit in $\ln P_{eq}$ is a factor of 2.7 in pressure. Whether that is negligible or disqualifying depends on the vessel, its heat exchange and its pressure limits, not on the regression model.

On the device side, heat and mass transfer in MH beds is well-described by continuum models that combine an energy balance, a gas balance, an equilibrium pressure law and a rate expression [8,11,26]. Machine learning has entered this side as well, through surrogate and physics-informed formulations [27–29] and data-driven identification of kinetic models [30]. The two lines of work are rarely connected. Tank models take thermodynamic inputs from measured pressure–composition–temperature (PCT) data, and composition models end with a predicted plateau pressure, which is not carried through to device performance. In addition, open tank-scale data sets are scarce, so neither side is usually tested against an independent device.

This work connects the two levels and uses that connection to ask a specific question: what accuracy in composition-predicted thermodynamics is sufficient for tank-level decisions? We (i) train composition-based models for ΔH and $\ln P_{eq}$ on ML-HydPARK, using a van 't Hoff-consistent parameterization and class-resolved error estimates; (ii) build a reduced-order, physics-based model of a cylindrical MH tank and calibrate it against the pressure, temperature and flow time series of a recently published open data set for a 7.8 kg AB_2 tank [31], using one charge rate for calibration and two for validation; (iii) use the composition-model prediction as a prior estimate of the plateau pressure and formation enthalpy of the tank alloy, and compare it with the values obtained by calibrating the tank model to the measured pressure, temperature and flow time series; and (iv) propagate the material-level error to tank performance through scenario analysis, variance-based sensitivity indices and an explicit tolerance window, before applying the calibrated tank to a virtual screen of the La–Ce–Ni and La–Ni–Al AB_5 series. The emphasis is on what the combination can and cannot establish from routinely logged tank data, rather than on a new regression model.

2. Materials and Methods

2.1. Thermodynamic Data

The material-level data were taken from ML-HydPARK v0.0.5 [19], a cleaned, machine-learning-ready version of the historical HydPARK database. It contains 420 entries with composition, formation enthalpy ΔH and entropy ΔS (reported as positive magnitudes per mol H_2), gravimetric capacity and the plateau pressure at 298 K expressed as $\ln(P_{eq}/P^0)$,

where $P^0 = 1$ atm is the standard pressure. Entries were grouped into the classes used in the source (AB_5, AB_2, AB, A_2B, Mg-based, multicomponent intermetallics (MIC) and solid solutions (SS)); the 23 entries added in later database versions without a class label were assigned by stoichiometry, giving 131 AB_2, 112 AB_5, 52 MIC, 49 SS, 34 AB, 32 Mg-based and 10 A_2B entries. Each entry was checked for internal consistency through the van 't Hoff relation; the largest deviation between reported and recomputed ln P_{eq} was 0.07, so no entry was removed. Four compositions occur twice after normalization with independent measurements and were retained.

2.2. Composition Descriptors and Equilibrium Models

Each composition was parsed with pymatgen [32] and described by 146 composition-based features from matminer [33]: the 132 Magpie elemental-property statistics [34], six stoichiometric norms and eight valence-orbital fractions. Gradient-boosting regression (GBR) [35] was used with the hyperparameters adopted by Witman et al. for hydride thermodynamics [22] (1500 trees, learning rate 0.005, depth 4, subsampling 0.75); random forests (500 trees) served as a baseline. Both were implemented in scikit-learn [36].

Further two parameterizations of the equilibrium were compared. In the independent scheme, ΔH and ΔS are regressed separately. In the van 't Hoff-consistent scheme, ΔH and ln P_{eq} at $T_0 = 298.15$ K are regressed, and the entropy follows from

$$\Delta S = R \ln(P_{eq,0}/P^0) + \Delta H/T_0, \quad (1)$$

where R is the universal gas constant and $P_{eq,0}$ is the plateau pressure at T_0, so that the predicted isotherm is anchored where the data are densest and only its temperature slope depends on ΔH. Because ΔH and ΔS are strongly correlated (enthalpy–entropy compensation), small independent errors in both can produce large errors in P_{eq}; the second scheme avoids this. Predictions were scored by the mean absolute error (MAE) in ln P_{eq} at 298 and 348 K, i.e., across the temperature range relevant for the tank studied below.

Accuracy was estimated by 5-fold cross-validation repeated three times with different splits, and, to probe extrapolation, by leave-class-out tests in which all AB_2 or all AB_5 entries were withheld from training. Out-of-fold residual pairs (e_H, e_{lnP}) were stored per class. Their empirical distribution, rather than the spread of a model ensemble, was used as the error model of a prediction, since ensembles trained on the same data underestimate the error for compositions unlike the training set.

2.3. Tank Data Set

The device-level data come from an open data set describing a prototype AB_2 tank designed to supply a 300 W proton-exchange-membrane fuel cell [31]. The vessel is a 316 L stainless-steel cylinder (3.5-inch schedule 80, wall 8.08 mm, 321 mm long, 2.53 L internal volume) holding 7.8 kg of $Ti_{0.815}Zr_{0.185}Cr_{0.77}Fe_{0.07}Mn_{0.88}$ with a nominal capacity of about 1.76 wt%. Flow, pressure and five shell thermocouples (TC1–TC5) were logged at 1 s intervals during three tests. In each test the activated bed was evacuated and charged with the mass-flow controller set to 5, 10 or 15 standard litres per minute (SLPM); the set flow was held only until the pressure reached the supply limit of about 24–25 bar (gauge), after 1.2, 0.44 and 0.29 h respectively, and charging then continued at the supply pressure for 18–28 h. The tank was subsequently discharged at a set flow of 5 SLPM until the pressure reached about 0.3–0.5 bar (gauge), after which the flow decayed to zero over several hours.

We consider that two features of the logs matter for modelling. First, the pressure readings are interpreted as gauge values (they read zero under vacuum and approach zero at the end of discharge to atmosphere), so 1.013 bar was added throughout. Second, during the slow, pressure-limited part of each charge the logged volumetric flow rate

alternates between zero and 0.5–1 SLPM, i.e., below the 1 SLPM lower limit of the controller. Integrating the readings as logged reproduces the absorbed masses reported with the data set; treating the zero readings as dropouts and interpolating across them gives 7–21% more (Table 1). The true absorbed mass was therefore taken to lie between these two bounds. Ambient temperature was estimated as the mean shell temperature at the end of each desorption log (25.3–26.2 °C), and the measured supply and cut-off pressures, initial bed temperatures and charge durations of each test were used as boundary conditions.

Table 1. Boundary conditions and observables of the three tests, derived from the time series of the open data set [31]. Pressures are absolute; temperatures are means of the five shell thermocouples, with single-thermocouple extremes in parentheses.

Quantity	5 SLPM	10 SLPM	15 SLPM
Role in this work	validation	calibration	validation
Initial/ambient temperature (°C)	26.1/25.3	29.8/26.2	23.8/26.1
Set-flow phase/total charge (h)	1.19/18.5	0.44/18.2	0.29/28.1
Supply pressure (bar)	25.4	25.0	24.7
H_2 absorbed, logged–dropout-corrected (g)	78.0–90.3	92.7–99.0	91.3–110.5
Peak shell temperature, charge (°C)	68.5 (70.2)	72.4 (74.2)	74.1 (75.5)
Rest before discharge (h)	5.6	0	0
H_2 released at 5 SLPM (g)	65.2	61.1	65.8
Minimum shell temperature, discharge (°C)	−11.4 (−18.0)	−11.3 (−14.5)	−12.7 (−17.0)
Cut-off pressure (bar)	1.51	1.31	1.51

2.4. Reduced-Order Tank Model

The bed was represented as a cylinder of inner radius $r_i = 42.7$ mm and effective length $L = V/(\pi r_i^2) = 0.442$ m, so that the model volume equals the reported internal volume; axial gradients and end effects were neglected. With an alloy density of 6.1 g cm^{-3} the bed porosity is $\varepsilon = 0.495$, and the theoretical capacity corresponds to $n_{max} = 68.1$ mol H_2 (137 g). The bed was divided into six annular finite volumes, each with a reacted fraction x_j and temperature T_j, surrounded by a lumped steel shell of temperature T_w, which is compared with the mean thermocouple reading. The equilibrium pressure combines the van 't Hoff term with a sloping plateau, a two-phase boundary term and hysteresis [26]:

$$\ln P_{eq}^{a,d} = \ln P_{eq,0} + (\Delta H/R)(1/T_0 - 1/T) + \beta(x - 0.5) + \zeta \ln[x/(1-x)] \pm \eta/2, \quad (2)$$

where x is the reacted fraction, T is the bed temperature, β is the plateau slope, ζ controls the width of the solid-solution and hydride single-phase regions, and $\eta = \ln(P^a/P^d)$ is the hysteresis factor (superscripts a and d denote absorption and desorption). The reaction rate follows an Avrami-type law with a logarithmic pressure driving force [8,37]:

$$dx/dt = k_m C_a \exp(-E_a/RT) \ln(P/P_{eq}^{a}) (1 - x) \text{ for } P > P_{eq}^{a}, \quad (3)$$

$$dx/dt = -k_m C_d \exp(-E_d/RT) \ln(P_{eq}^{d}/P) x \text{ for } P < P_{eq}^{d}, \quad (4)$$

where P is the gas pressure, with the pre-exponential factors and activation energies of Jemni and Ben Nasrallah ($C_a = 59.2$ s^{-1}, $E_a = 21.18$ kJ mol^{-1}, $C_d = 9.57$ s^{-1}, $E_d = 16.45$ kJ mol^{-1}) [8] scaled by a single multiplier k_m. The energy balance of volume j is

$$m_j c_p \, dT_j/dt = \Sigma \, G_{j,k}(T_k - T_j) + n_{max,j} \, \Delta H \, dx_j/dt, \quad (5)$$

where m_j is the alloy mass of volume j, $c_p = 480$ J kg^{-1} K^{-1} is the specific heat of the alloy, $G_{j,k}$ are radial conductances computed from an effective bed conductivity λ_{eff}, and the outermost volume exchanges heat with the shell through a contact coefficient h_{bw}. The shell (10 kg, 500 J kg^{-1} K^{-1}) loses heat to ambient air at T_{amb} through an external coefficient h_{out}

over 0.157 m^2. Hydrogen in the void volume (εV plus 0.1 L of piping) is treated as an ideal gas at the mean bed temperature:

$$dn_g/dt = \dot{n}_{in} - \Sigma_j\, n_{max,j}\, dx_j/dt,\ P = n_g R T^{-}/V_g. \quad (6)$$

Here n_g is the amount of hydrogen in the void volume V_g, $\dot{n}_{in}$ is the net molar flow of hydrogen into the tank (negative during discharge) and T^{-} is the mean bed temperature. The test protocol is reproduced as logged. Charging starts from an evacuated bed at the measured initial temperature; the inflow equals the set flow until P reaches the measured supply pressure P_{sup}, after which the pressure is held and the flow is set by absorption, for the measured charge duration. Where the log shows a rest, the closed tank relaxes towards T_{amb}. Discharge draws 5 SLPM until P falls to the measured cut-off P_{cut}, after which the pressure is held and the outflow is set by desorption. The stiff system of ordinary differential equations was integrated with a backward-differentiation solver (SciPy), and simulating one complete test takes about 0.3 s. This short run time, rather than that of a two- or three-dimensional solution, is what makes the analyses below practical: the calibration requires of the order of 10^4 simulations of complete tests, and the scenario, tolerance-window, screening and Sobol analyses a further ~3000 simulations of the design protocol.

2.5. Calibration and Validation

Nine parameters were estimated (Table 2): ln $P_{eq,0}$ and ΔH of the tank alloy, β, ζ, η, λ_{eff}, h_{bw}, h_{out} and ln k_m. They were fitted to the 10 SLPM test by bounded weighted least squares (trust-region reflective algorithm) from four starting points, one of which was the composition-model prediction. Six residual blocks entered the objective, each residual being divided by the standard deviation σ assumed for the corresponding measurement, which sets its weight: shell temperature during charging (5 min resolution, $\sigma = 2$ °C), pressure during charging ($\sigma = 0.7$ bar), absorbed hydrogen every 30 min, penalized only outside the logged–dropout-corrected band ($\sigma = 3$ g), so that the dropout uncertainty enters the objective as an interval rather than as a point value, and shell temperature, pressure and delivered hydrogen during the first 6, 3 and 6 h of discharge (2 min resolution; $\sigma = 2$ °C, 0.5 bar and 2 g). Each block was scaled so that blocks of different length carry equal weight. The 5 and 15 SLPM tests were not used for calibration and serve as validation. Because the time-series residuals are strongly autocorrelated, formal confidence intervals from the Jacobian would be far too narrow.

Table 2. Parameters of the tank model: bounds used in calibration, estimates from the 10 SLPM test, the range over the 12 solutions fitted to the three tests, and the ranges used for the Sobol analysis.

Parameter	Bounds	10 SLPM Fit	Range, 3 Tests	Sobol Range
ln($P_{eq,0}$/atm)	−1.5 to 4.0	2.68	2.59–2.68	±1.75
ΔH (kJ mol^{-1})	12 to 45	28.9	27.2–29.6	±7
β	0 to 2	0.50	0.04–0.52	0 to 1
ζ	0.03 to 0.40	0.40 [a]	0.32–0.40	fixed
η	0 to 1.2	0.31	0.08–0.49	0 to 0.6
λ_{eff} (W m^{-1} K^{-1})	0.3 to 5	5.0 [a]	3.4–5.0	0.5 to 5
h_{bw} (W m^{-2} K^{-1})	100 to 3000	3000 [a]	2020–3000	100 to 3000
h_{out} (W m^{-2} K^{-1})	2 to 40	6.8	5.2–7.0	3 to 20
k_m	0.05 to 20	0.06	0.05–1.1	0.05 to 20

[a] At the upper bound; see Section 4.3.

Parameter uncertainty was therefore assessed empirically, by refitting the model separately to each of the three tests from the same starting points and retaining all solutions within 10% of the best objective (12 parameter sets), and by a profile of the objective over

ln $P_{eq,0}$, in which all other parameters were re-optimized at fixed plateau values. The calibrated plateau was then compared with the composition-model prediction and its AB_2 error distribution, in the spirit of the calibration of computer models against observations [38].

2.6. Propagating Material Uncertainty to the Tank

All scenarios use a common design protocol based on the test rig: an evacuated tank at 26 °C is charged with a set flow of 10 SLPM and a 25.4 bar supply limit for 18 h, then discharged at 5 SLPM down to 1.4 bar for 7 h. Four key performance indicators (KPIs) were recorded: the hydrogen absorbed in 18 h, the hydrogen delivered, the minimum shell temperature during discharge and the time needed to absorb 60 g.

Scenarios: for each of the 12 calibrated parameter sets, the transport parameters were kept and the thermodynamics replaced by (A) the calibrated values, (B) the two nearest AB_2 analogues in the database, $Ti_{0.8}Zr_{0.2}CrMn$ and $Zr_{0.2}Ti_{0.8}MnCr$, (C_0) the GBR point prediction and (C) the GBR prediction perturbed by nine randomly drawn AB_2 out-of-fold residual pairs.

Sensitivity: first-order and total Sobol indices [39,40] were computed with SALib [41] for eight parameters, with ln $P_{eq,0}$ and ΔH varied by ± 1.75 and ± 7 kJ mol^{-1} around the calibrated values (the AB_2 cross-validation MAEs) and β, η, λ_{eff}, h_{out}, k_m and h_{bw} over wide engineering ranges (See Table 2).

Tolerance window and screening: ln $P_{eq,0}$ and ΔH were varied on a grid around the calibrated values to map the region in which the KPIs stay within 10% of their calibrated values, and the calibrated tank was used to compare AB_5 candidates from the $La_{1-x}Ce_xNi_5$ and $LaNi_{5-x}Al_x$ series using predicted and, where available, database thermodynamics.

3. Results

Figure 1 presented below summarizes the workflow. The composition model supplies a prediction, with an error model, for the equilibrium of the tank alloy, the tank model maps equilibrium and transport parameters to the logged pressure, temperature and flow, and the open tank data calibrate and validate it. The same tank model is then used as a common yardstick for the thermodynamic scenarios, the sensitivity analysis and the AB_5 screen.

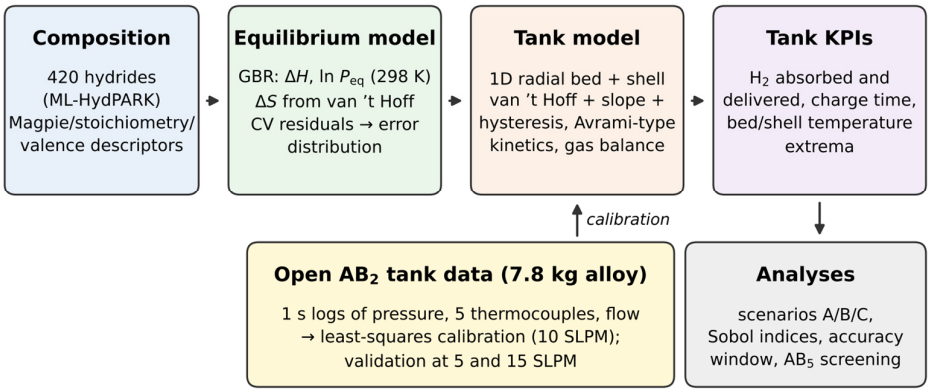


Figure 1. Workflow linking composition, equilibrium thermodynamics and tank performance. GBR, gradient-boosting regression. KPI, key performance indicator. SLPM, standard litres per minute.

3.1. Accuracy of Composition-Based Thermodynamics

Our Table 3 compares the two parameterizations. Over all 420 entries, the GBR model reaches an MAE of 7.2 kJ mol^{-1} in ΔH ($R^2 = 0.62$) and of 2.03 in ln P_{eq} at 298 K. Restricted to the 345 entries with $|\ln P_{eq}| < 6$, i.e., plateaus between about 0.0025 and 410 bar, the MAE falls to 1.33 at 298 K and 1.19 at 348 K. The van 't Hoff-consistent scheme is better than the

independent scheme in every case, but only by 0.04–0.05 in ln P_{eq}; the gain is consistent across repeated splits and both learners, yet it is small compared with the error itself. Random forests are slightly worse than GBR. As in earlier work [22], the most important descriptors for both targets are the mean and mode of the elemental ground-state volume per atom, followed by covalent radius and d-shell filling, which reflects the interstitial-volume control of hydride stability.

Table 3. Cross-validated accuracy of the equilibrium models (5-fold, three repeats; mean over repeats). The window subset contains the 345 entries with $|\ln P_{eq,0}| < 6$.

Model	Scheme	MAE ΔH (kJ mol^{-1})	MAE ln P_{eq}, 298 K, All	MAE ln P_{eq}, 298 K, Window	MAE ln P_{eq}, 348 K, Window
GBR	independent	7.19	2.07	1.38	1.25
GBR	van 't Hoff	7.19	2.03	1.33	1.19
RF	independent	7.42	2.18	1.43	1.27
RF	van 't Hoff	7.42	2.11	1.35	1.20

Standard deviations across repeats are ≤0.06 for all ln P_{eq} entries and ≤0.22 kJ mol^{-1} for ΔH.

Accuracy depends strongly on the alloy class (Figure 2c, Table 4). AB_5 compounds, the most homogeneous and best-sampled class, are predicted to within 0.90 in ln P_{eq} and 3.3 kJ mol^{-1} in ΔH. AB_2 Laves phases, which are more numerous but chemically diverse, have twice the error (1.75 and 7.0 kJ mol^{-1}), and Mg-based, AB and multicomponent compounds are worse still. When an entire class is withheld, the error grows to 2.12 for AB_5 and to 3.89 for AB_2, with a systematic underprediction of the plateau (mean signed error −1.42 and −3.23, respectively). For the AB_2 out-of-fold residuals, the standard deviations are 3.05 in ln P_{eq} and 10.3 kJ mol^{-1} in ΔH, with a correlation of −0.60 between the two errors: this distribution is the error model carried forward to the tank.

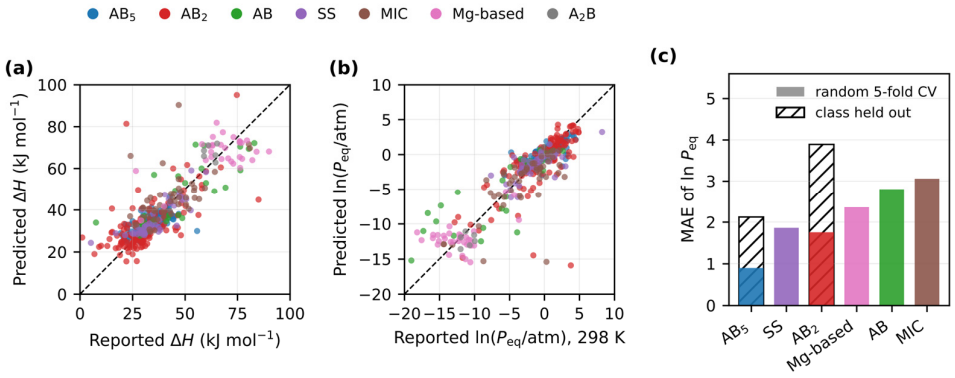


Figure 2. Composition-based equilibrium models. (**a**) Out-of-fold predicted versus reported ΔH; (**b**) out-of-fold predicted versus reported ln P_{eq} at 298 K, coloured by alloy class; (**c**) mean absolute error of ln P_{eq} by class for random cross-validation (filled) and with the class withheld from training (hatched).

For the tank alloy $Ti_{0.815}Zr_{0.185}Cr_{0.77}Fe_{0.07}Mn_{0.88}$, which is not in the database, the model predicts ΔH = 26.7 kJ mol^{-1} and ln $P_{eq,0}$ = 0.63, i.e., a plateau of about 1.9 bar at 298 K and ΔS = 94.9 J mol^{-1} K^{-1}. The two nearest database analogues bracket this prediction: $Ti_{0.8}Zr_{0.2}CrMn$ (ΔH = 24.6 kJ mol^{-1}, ln $P_{eq,0}$ = 3.17) and $Zr_{0.2}Ti_{0.8}MnCr$ (30.0 kJ mol^{-1}, 1.80). The spread between these two related compositions, 1.4 units in ln $P_{eq,0}$, is itself comparable to the class error.

Table 4. Class-resolved out-of-fold errors of the GBR model (van 't Hoff scheme) and leave-class-out errors.

Class	n	MAE ΔH (kJ mol^{-1})	MAE ln P_{eq}	Leave-Class-Out MAE ln P_{eq} (Bias)
AB_5	112	3.3	0.90	2.12 (−1.42)
AB_2	131	7.0	1.75	3.89 (−3.23)
SS	49	5.5	1.86	–
Mg-based	32	9.9	2.35	–
AB	34	8.9	2.80	–
MIC	52	9.2	3.06	–

3.2. Calibration and Validation of the Tank Model

All four starting points of the 10 SLPM fit converged to nearly the same solution (objective 153.8–155.8), including the start at the composition-model prediction (ln $P_{eq,0} = 0.63$). The estimates are listed in Table 2. The plateau, ln $P_{eq,0} = 2.68$, and $\Delta H = 28.9$ kJ mol^{-1} are well-defined; together they give $\Delta S \approx 119$ J mol^{-1} K^{-1}, a typical value for AB_2 hydrides. The tank also fixes a moderately sloping plateau ($\beta \approx 0.5$), hysteresis $\eta \approx 0.3$ and an external heat-transfer coefficient of 6.8 W m^{-2} K^{-1}, consistent with natural convection and radiation from a bare cylinder in still air. Bed conductivity and bed–wall contact move to their upper bounds, i.e., the data are best described with negligible internal resistance (at the fitted values the external resistance, $1/(h_{out} A_{out}) \approx 0.94$ K W^{-1}, is 20–30 times the combined bed and contact resistance, so the objective is nearly flat in λ_{eff} and h_{bw}), and the kinetic multiplier is not identified (0.05–1.1 across fits): with a supply-limited charge and a slow discharge, the tank operates close to equilibrium and its heat balance, not its kinetics, sets the rates.

For the 10 SLPM calibration test, the root-mean-square errors are 3.3 °C for shell temperature and 0.4 bar for pressure during charging, and 2.6 °C, 0.6 bar and 1.7 g for temperature, pressure and delivered hydrogen during the first hours of discharge. Figure 3 compares the calibrated model with all three tests. Without refitting, the 5 and 15 SLPM tests are reproduced with errors of 5.5 and 3.6 °C (charging temperature), 1.0 and 0.4 bar (charging pressure), 4.4 and 3.5 °C (discharge temperature), 0.9 and 0.7 bar (discharge pressure) and 2.6 and 1.8 g (delivered hydrogen). The minimum mean shell temperature during discharge is predicted within 1.3 °C in all tests (−10.0 to −11.1 °C against −9.6 to −12.4 °C at 5 min resolution).

Even so, two systematic deviations remain. The peak shell temperature during the set-flow phase is underpredicted by 11–15 °C, because the bed equilibrium temperature at the supply pressure is set by the low-concentration part of the isotherm, which the simple form of Equation (2) represents poorly; the long cooling tail that controls the absorbed amount is, however, reproduced. The model also predicts a concave absorption curve, whereas the logged uptake stalls between about 4 and 10 h and then accelerates. The stall coincides with the period in which the flow readings fall below the controller's rated range, and the corrected band in Figure 3 shows that much of it is a metering artefact, but the late acceleration is not captured.

3.3. What the Tank Implies for the Composition Model

The plateau is the best-determined parameter of the tank. Refitting to each test separately gives ln $P_{eq,0} = 2.67$ (5 SLPM), 2.68 (10 SLPM) and 2.59 (15 SLPM), and the 12 solutions within 10% of the best objective span 2.59–2.68, i.e., 13.5–14.8 bar at 298 K; ΔH spans 27.2–29.6 kJ mol^{-1} (Figure 4b). The profile in Figure 4a rises steeply on both sides of its minimum: fixing the plateau 0.25 units too low or too high raises the objective by 69 and 20 units, respectively, and at the GBR prediction the objective is about 35 times its

minimum. Compared with these estimates, the GBR model predicts ΔH almost correctly (26.7 kJ mol^{-1}) but places the plateau 2.0 units too low, a factor of about 7 in pressure, the error lies almost entirely in the entropy term. The direction of the error matches the systematic underprediction found when AB_2 alloys were withheld (See Table 4), and its size is within the 1σ ellipse of the AB_2 error model. In other words, the prediction is wrong by an amount the model itself declares plausible, and three ordinary tank tests constrain the plateau roughly two orders of magnitude more tightly than the composition model does.

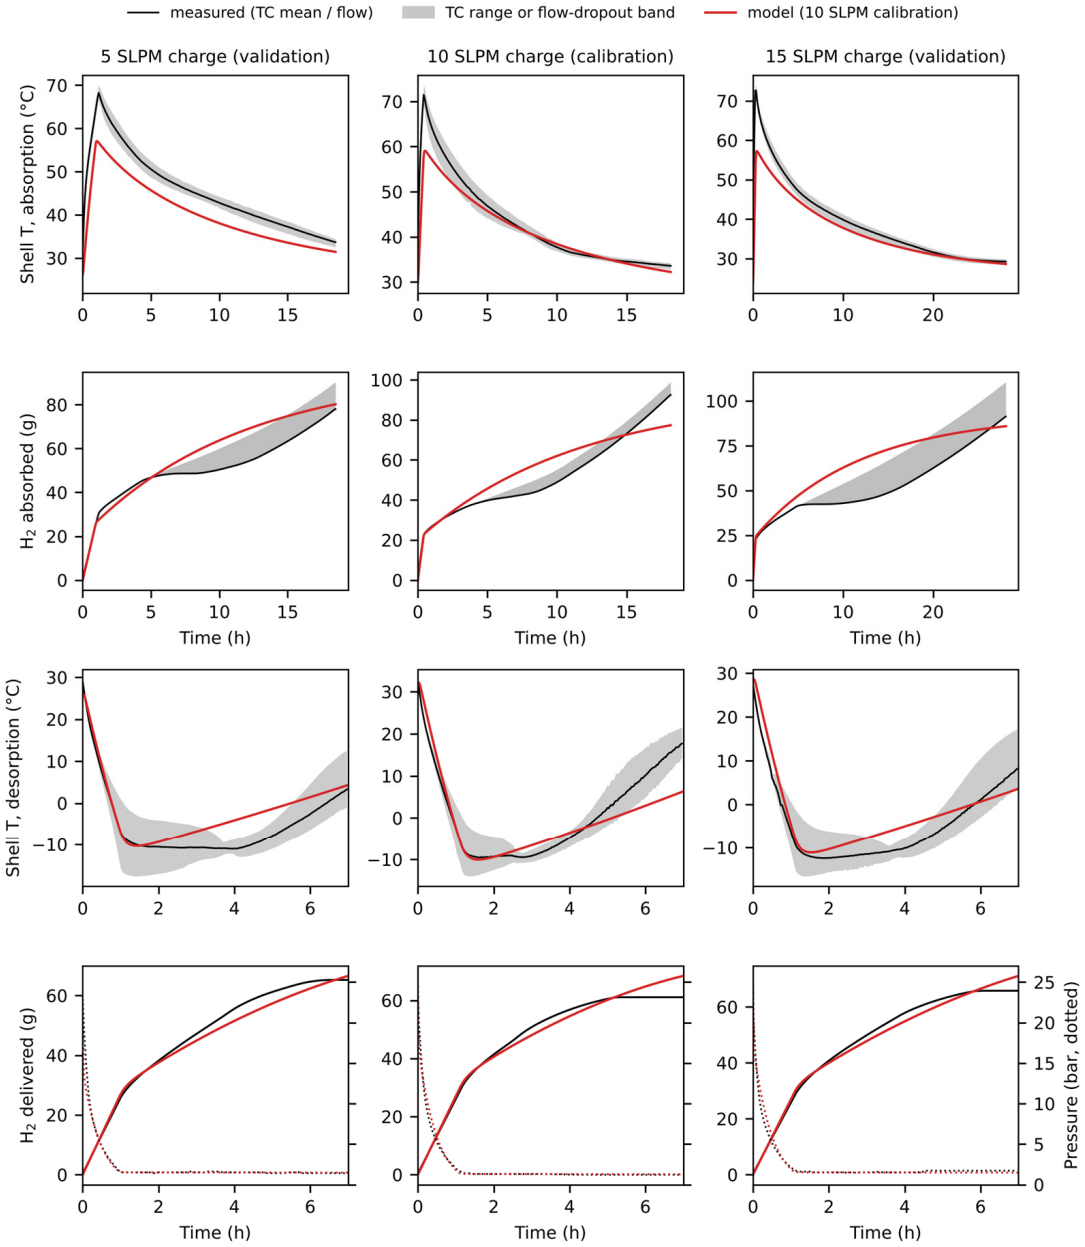


Figure 3. Calibrated tank model (red) against the three tests (black). Rows: mean shell temperature during charging; hydrogen absorbed, with the band between logged and dropout-corrected flow; mean shell temperature during discharge, with the range of the five thermocouples; hydrogen delivered (solid) and absolute pressure (dotted) during discharge. Parameters were fitted to the 10 SLPM test only.

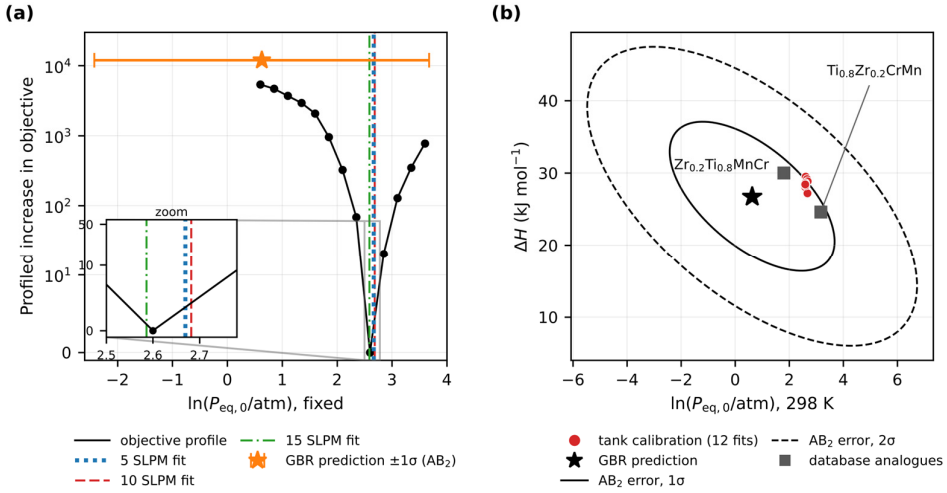


Figure 4. Plateau of the tank alloy. (**a**) Profile of the calibration objective over ln $P_{eq,0}$, black line and markers (all other parameters re-optimized, symmetric-log scale), with the separate fits to the three tests (vertical lines; the inset enlarges the region around the minimum) and the GBR prediction with its AB_2 1σ interval (one standard deviation, σ, of the AB_2 out-of-fold residuals). (**b**) Calibrated solutions (red) in the (ln $P_{eq,0}$, ΔH) plane, the GBR prediction with its 1σ (solid) and 2σ (dashed) AB_2 error ellipses, and the two nearest database analogues (grey squares).

3.4. Propagation of Material-Level Uncertainty to the Tank

Table 5 and next Figure 5 compare the design protocol for the 12 calibrated transport parameter sets under different thermodynamic inputs. With the calibrated thermodynamics (A), the tank absorbs 79 g in 18 h, needs 8.8 h to take up 60 g, peaks at 58 °C and delivers 69 g at 5 SLPM, its shell falling to −11 °C. The nearest database analogues behave quite differently in the same vessel. $Ti_{0.8}Zr_{0.2}CrMn$ (B1), with a plateau 0.5 units higher, is held back by the supply pressure and absorbs only 55 g, never reaching 60 g. $Zr_{0.2}Ti_{0.8}MnCr$ (B2), 0.8 units lower, absorbs 118 g and reaches 60 g in 3.8 h, and because it stores more it delivers almost as much as A (65 g). The GBR point prediction (C_0) behaves like a much more stable hydride: the tank takes up 60 g in 1.1 h and reaches its theoretical capacity during the charge, the shell reaches 106 °C, and only 47 g of the 139 g stored are delivered before the cold bed reaches the cut-off pressure, 31% less than with the calibrated values. When the GBR prediction is perturbed by AB_2 cross-validation errors (C), the 5–95% interval of delivered hydrogen spans 3 to 106 g, and only 6 of 108 draws fall within ±10% of the calibrated delivery.

Table 5. Tank KPIs for the design protocol under different thermodynamic inputs: median (5–95% interval) over the calibrated transport parameter sets.

Scenario	H_2 Absorbed in 18 h (g)	Time to 60 g (h)	Peak Shell T (°C)	H_2 Delivered (g)
A: calibrated	79 (76–82)	8.8 (8.2–10.2)	58 (56–62)	69 (68–70)
B1: $Ti_{0.8}Zr_{0.2}CrMn$	55 (46–59)	not reached	49 (45–52)	54 (46–58)
B2: $Zr_{0.2}Ti_{0.8}MnCr$	118 (116–120)	3.8 (3.4–4.2)	75 (72–78)	65 (64–68)
C_0: GBR point	139 (139–140)	1.1	106 (105–111)	47 (42–49)
C: GBR + CV error	134 (16–140)	1.1 (1.1–4.6) [a]	93 (35–143)	36 (3–106)

Thermodynamics (ΔH in kJ mol^{-1}; ln $P_{eq,0}$): A 28.6; 2.62 (ensemble mean), B1 24.6; 3.17, B2 30.0; 1.80, C_0 26.7; 0.63. [a] Over the draws that reach 60 g.

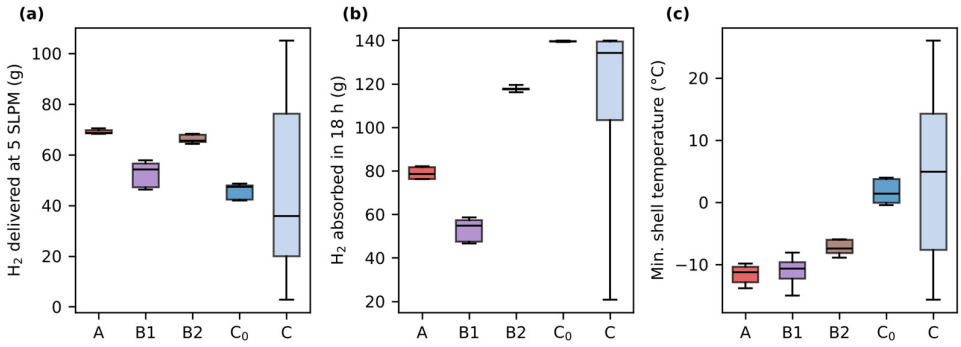


Figure 5. Distributions of (**a**) hydrogen delivered at 5 SLPM, (**b**) hydrogen absorbed in an 18 h charge and (**c**) minimum shell temperature during discharge under the thermodynamic scenarios of Table 5. Boxes show quartiles and whiskers the 5–95% range.

3.5. Tolerance Window and Sensitivity

Delivery is a non-monotonic function of the plateau. Figure 6a maps the delivered hydrogen over offsets of ln $P_{eq,0}$ and ΔH from the calibrated values. It peaks (73 g) for a plateau 0.5 units lower than calibrated, because a more stable hydride stores more during the supply-limited charge, and it falls steeply for higher plateaus, which leave the tank partly empty. Delivery stays within 10% of its calibrated value (68.2 g) for offsets from −1.46 to +0.24, a window of 1.71 units that is strongly asymmetric: an overestimated plateau is far more costly than an underestimated one. The charging behaviour is much less tolerant. The hydrogen absorbed in 18 h stays within ±10% only for plateau offsets between −0.15 and +0.15, and an offset of ±0.25 changes the time to absorb 60 g from 8.7 h to 6.2 and 13.8 h. Along the ΔH axis at fixed plateau, delivery stays within 10% for any lower enthalpy and up to +3.1 kJ mol^{-1} higher, and the absorbed amount for offsets between −5.5 and +5.0 kJ mol^{-1}. For comparison, the half-width of the delivery window (0.85) is about equal to the AB_5 cross-validation MAE (0.90) and half the AB_2 MAE (1.75), whereas the absorption window is six and twelve times narrower than these errors.

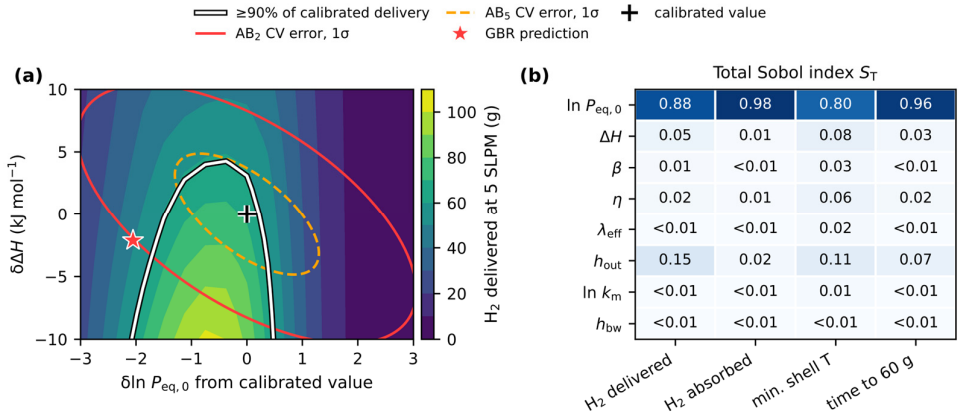


Figure 6. (**a**) Hydrogen delivered at 5 SLPM under the design protocol as a function of offsets in ln $P_{eq,0}$ and ΔH from the calibrated values, the white contour encloses delivery within 90% of the calibrated value, ellipses show the 1σ cross-validation errors of the AB_2 (red) and AB_5 (orange) models, and the star marks the GBR prediction for the tank alloy. (**b**) Total Sobol indices of four tank KPIs with thermodynamic uncertainty equal to the AB_2 MAEs and transport parameters varied over the ranges of Table 2.

The Sobol indices in Figure 6b express the same point globally. With the thermodynamics varied by ±1.75 in ln $P_{eq,0}$ and ±7 kJ mol^{-1} in ΔH (the AB_2 MAEs) and the transport

parameters over the ranges in Table 2, the plateau alone carries total indices of 0.98 for absorbed hydrogen, 0.96 for the time to absorb 60 g, 0.88 for delivered hydrogen and 0.80 for the minimum shell temperature. The external coefficient h_{out} is the only other parameter of consequence (0.15 for delivery and 0.11 for minimum temperature); ΔH contributes at most 0.08, and β, η, λ_{eff}, k_m and h_{bw} at most 0.06. The 95% bootstrap confidence intervals are at most ± 0.15 (2560 model runs, none failed; runs that did not take up 60 g within 18 h, 38% of the total, were assigned the longest time observed).

3.6. Virtual Screening of AB_5 Substitution Series

To illustrate how the calibrated tank can be used despite these limits, the AB_2 charge was replaced by the same packed volume of AB_5 alloy (10.8 kg, 1.4 wt%; $\beta = 0.25$, $\eta = 0.25$ and $LaNi_5$ kinetics [8]), keeping the vessel, heat exchange and design protocol fixed. The GBR model predicts plateaus at 298 K that rise from 1.5 bar for $LaNi_5$ to 3.3, 3.8, 3.9 and 6.9 bar for $La_{1-x}Ce_xNi_5$ with $x = 0.1$, 0.2, 0.25 and 0.5, and fall to 1.1, 0.83, 0.50, 0.37, 0.26 and 0.09 bar for $LaNi_{5-x}Al_x$ with $x = 0.1$, 0.2, 0.3, 0.4, 0.55 and 0.75. With predicted thermodynamics every candidate fills the tank to about 154 g in the 18 h charge, but the hydrogen delivered at 5 SLPM without heating follows the plateau monotonically (See Figure 7): 20 g for $LaNi_5$, 59–69 g for $x_{Ce} = 0.1$–0.25, 102 g for $La_{0.5}Ce_{0.5}Ni_5$, and 3–10 g for the Al-substituted alloys. Where database thermodynamics exist they change the absolute numbers (50 g instead of 20 g for $LaNi_5$, whose database plateau is 0.7 units higher) but not the ranking. With the 25 bar supply the AB_5 beds would reach 67–173 °C during charging, so a real AB_5 design would also lower the supply pressure or add cooling.

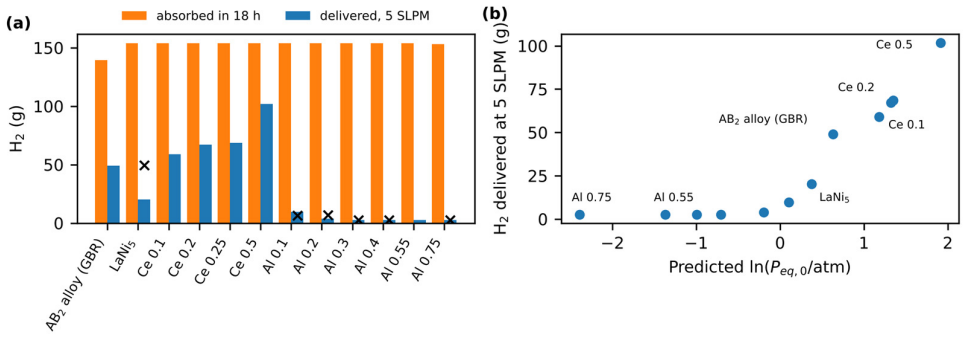


Figure 7. Virtual screening of AB_5 alloys in the calibrated tank under the design protocol. (**a**) Hydrogen absorbed in 18 h and delivered at 5 SLPM with GBR-predicted thermodynamics, where crosses show delivery with database thermodynamics where available. (**b**) Delivered hydrogen versus predicted ln $P_{eq,0}$. The AB_2 point uses the GBR prediction for the tank alloy, not the calibrated value.

4. Discussion

4.1. How Accurate Must Predicted Thermodynamics Be?

For this tank the answer depends on which decision the prediction is meant to support. If the question is how much hydrogen a full charge will deliver without heating, the room-temperature plateau must be known to within roughly $-1.5/+0.25$ in ln P_{eq}; the window is asymmetric because an overestimated plateau leaves the tank under-filled, while a moderately underestimated one is compensated by a larger stored amount. If the question is how much hydrogen an 18 h charge stores, how long it takes or how hot the bed becomes, the plateau must be known to within about ± 0.15, i.e., $\pm 15\%$ in pressure. The best class-resolved accuracy available from composition alone is 0.90 for AB_5 and 1.75 for AB_2, and nearly 4 when the composition lies outside the well-sampled part of the class. Composition models are therefore close to adequate for a first estimate of deliverable capacity of AB_5 alloys, but

they are six to twelve times too coarse for predicting how a supply-limited tank charges. This gap cannot be closed by better transport modelling: once the plateau is uncertain at this level it dominates every KPI (Figure 6b). The specific numbers depend on the vessel, its heat exchange and its pressure limits. The method for obtaining them, a calibrated tank model plus a tolerance map, transfers directly to other vessels.

The tank alloy also illustrates where the composition-model error comes from. The predicted enthalpy is within 2 kJ mol^{-1} of the calibrated value, while the plateau is off by 2.0 units; the error therefore sits in the entropy, which composition descriptors capture poorly and which the database reports with considerable scatter between nominally similar alloys. $Ti_{0.815}Zr_{0.185}Cr_{0.77}Fe_{0.07}Mn_{0.88}$ is A-rich (B/A ≈ 1.72) and Fe-substituted, and the closest database entries differ from each other by 1.4 units in ln $P_{eq,0}$. A model trained on such data cannot be expected to resolve plateau differences of a few tenths for this class. Reporting class-resolved and leave-class-out errors, as done here, at least makes this visible to the user.

None of this makes composition models useless for system design. Their errors are large in absolute terms but largely systematic within a substitution series, so rankings survive: in Figure 7 the predicted order of the La–Ce–Ni and La–Ni–Al alloys is unchanged when database values replace predictions. The screen also reproduces the design logic of systems built around these alloys. Ce substitution raises the plateau and is the route to unheated delivery at fuel-cell pressures [16]. Al substitution lowers it, which is useful for absorbing near-atmospheric hydrogen from solid-oxide electrolysers but requires heating to 90–120 °C for release [18]; in an unheated tank these alloys deliver almost nothing. A practical division of labour follows: use composition models to choose the substitution direction and narrow a series to two or three candidates, then fix ln P_{eq} with one measured isotherm before committing to a vessel design. Section 3.3 shows that a few well-logged tank tests interpreted through a calibrated model can serve the same purpose.

4.2. Heat Transfer in Context

The calibrated tank operates close to equilibrium. Its charge is limited by the supply pressure and its discharge by the cut-off pressure, and in both cases the rate is set by how fast heat crosses the outer surface (h_{out} ≈ 6–7 W m^{-2} K^{-1}). Bed conductivity and bed–wall contact moved to the upper bounds of the fit, and the kinetic multiplier was not identified. This should not be read as a general statement that bed heat transfer or kinetics are unimportant, which would contradict a large body of work [10–12]. It means that in a bare vessel cooled by still air the external resistance dominates, so internal enhancements cannot change the outcome.

The same calibration shows where engineering effort would pay: the external coefficient is the only transport parameter with a substantial Sobol index, and a finned shell or water jacket would both shorten the charge and widen the plateau tolerance window. Thermal management and material selection can be traded against each other, but only once both are quantified on the same model.

4.3. Limitations

The calibration uses one tank and three tests. Within that scope, it identifies the plateau, enthalpy, plateau slope, hysteresis and external heat transfer, but not the internal transport parameters, two of which (λ_{eff}, h_{bw}) and the phase-boundary term ζ ended at their bounds. These should be read as effective values that absorb model-form error: the lumped shell is compared with the mean of five surface thermocouples, axial conduction and the end caps are neglected, and Equation (2) is a simple isotherm form. The latter is the most likely cause of the two systematic deviations in Figure 3: the 11–15 °C underprediction of the

peak shell temperature, which is set by the low-concentration part of the isotherm, and the missing late acceleration of the absorption curve. The thermocouples also show a plateau near −10 °C between about 1.5 and 4 h of each discharge that the model does not reproduce; frost formation on the cold shell would produce such a plateau, but the data do not allow this to be confirmed. The absorbed masses are uncertain by 7–21% because the controller operated below its rated range during most of each charge, and the sub-second dynamics of the controller are not modelled.

The ensemble of 12 fitted parameter sets is a pragmatic substitute for a full posterior. On the material side, ML-HydPARK contains few sub-stoichiometric Ti–Zr–Cr–Mn–Fe Laves phases, and its entries come from different measurement protocols, so part of the cross-validation error is data noise rather than model error. Finally, the AB_5 screen keeps the AB_2 vessel and supply pressure, under which AB_5 beds would overheat, and it is meant to rank alloys, not to design an AB_5 tank.

These two extensions follow directly. A more flexible isotherm, for example one derived from a phase-equilibrium model [26], fitted jointly to tank and PCT data, would remove most of the remaining systematic error. And because the framework expresses every material uncertainty in tank-level units, it can rank candidate experiments (a PCT isotherm, a second tank test, a different flow rate) by how much they would reduce the uncertainty of a design decision.

5. Conclusions

We linked composition-based predictions of hydride thermodynamics to a physics-based model of a 7.8 kg AB_2 tank, calibrated on the pressure, temperature and flow time series of one open test and validated on two others, and used the link to ask how accurate the predictions must be for tank-level decisions. Gradient-boosting models trained on 420 ML-HydPARK entries predict the room-temperature plateau with cross-validated MAEs of 0.90 (AB_5) and 1.75 (AB_2) in $\ln P_{eq}$. The calibrated tank model reproduces the validation tests within 3.5–5.5 °C and 0.4–1.0 bar and places the alloy's plateau at $\ln P_{eq,0} = 2.59$–2.68, 2.0 units above the composition-model prediction, while the predicted enthalpy is nearly correct. Substituting the predicted thermodynamics changes the simulated tank qualitatively: it takes up 60 g in about one hour instead of nine, overheats to 106 °C and delivers 31% less hydrogen. Delivered capacity tolerates plateau errors of about −1.5/+0.25, whereas charging behaviour tolerates only ±0.15, and the plateau dominates the variance of every KPI. Composition models are therefore useful for ranking substitution series, as the La–Ce–Ni and La–Ni–Al screen shows, but a measured isotherm, or well-logged tank tests interpreted through a calibrated model, remains necessary before tank performance can be predicted.

Author Contributions: Conceptualization, A.K. and D.C.; methodology, A.K. and F.N.; software, A.K. and F.N.; validation, A.K. and F.N.; formal analysis, A.K. and F.N.; investigation, A.K. and F.N.; data curation, F.N.; writing—original draft preparation, A.K.; writing—review and editing, A.K., F.N. and D.C.; visualization, A.K. and F.N.; supervision, D.C.; project administration, D.C. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: This study analyzed publicly available data. The hydride thermodynamics are from ML-HydPARK v0.0.5 (https://doi.org/10.5281/zenodo.10680097) [19], and the tank time series are from Mendeley Data (https://doi.org/10.17632/8km9z62tct.1) [31]. The code for the equilibrium models, the tank model, the calibration and the sensitivity analysis, together with the processed data and the figures, is openly available at https://github.com/akartbayev/mh-composition-to-tank (accessed on 16 November 2026).

Acknowledgments: The authors thank the creators of the ML-HydPARK database and of the open AB_2 tank data set for making their data publicly available. During the preparation of this manuscript, the authors used NotebookLM, Google Cloud Platform (GCP) and Gemini for the purposes of language editing, and Qwen3.8-27B and Visual Studio Code (v1.147.0) for the purposes of coding. The authors have reviewed and edited the output and take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

CV	Cross-validation
GBR	Gradient-boosting regression
KPI	Key performance indicator
MAE	Mean absolute error
MH	Metal hydride
MIC	Multicomponent intermetallic compound
PCT	Pressure–composition–temperature
RMSE	Root-mean-square error
RF	Random forest
SLPM	Standard litres per minute
SS	Solid solution
TC	Thermocouple

References

1. Schlapbach, L.; Züttel, A. Hydrogen-storage materials for mobile applications. *Nature* **2001**, *414*, 353–358. [CrossRef]
2. Sakintuna, B.; Lamari-Darkrim, F.; Hirscher, M. Metal hydride materials for solid hydrogen storage: A review. *Int. J. Hydrogen Energy* **2007**, *32*, 1121–1140. [CrossRef]
3. Hirscher, M.; Yartys, V.A.; Baricco, M.; Bellosta von Colbe, J.; Blanchard, D.; Bowman, R.C., Jr.; Broom, D.P.; Buckley, C.E.; Chang, F.; Chen, P.; et al. Materials for hydrogen-based energy storage—Past, recent progress and future outlook. *J. Alloys Compd.* **2020**, *827*, 153548. [CrossRef]
4. Klopčič, N.; Grimmer, I.; Winkler, F.; Sartory, M.; Trattner, A. A review on metal hydride materials for hydrogen storage. *J. Energy Storage* **2023**, *72*, 108456. [CrossRef]
5. Bellosta von Colbe, J.; Ares, J.-R.; Barale, J.; Baricco, M.; Buckley, C.; Capurso, G.; Gallandat, N.; Grant, D.M.; Guzik, M.N.; Jacob, I.; et al. Application of hydrides in hydrogen storage and compression: Achievements, outlook and perspectives. *Int. J. Hydrogen Energy* **2019**, *44*, 7780–7808. [CrossRef]
6. Lototskyy, M.V.; Yartys, V.A.; Pollet, B.G.; Bowman, R.C., Jr. Metal hydride hydrogen compressors: A review. *Int. J. Hydrogen Energy* **2014**, *39*, 5818–5851. [CrossRef]
7. Lototskyy, M.V.; Davids, M.W.; Tolj, I.; Klochko, Y.V.; Sekhar, B.S.; Chidziva, S.; Smith, F.; Swanepoel, D.; Pollet, B.G. Metal hydride systems for hydrogen storage and supply for stationary and automotive low temperature PEM fuel cell power modules. *Int. J. Hydrogen Energy* **2015**, *40*, 11491–11497. [CrossRef]
8. Jemni, A.; Ben Nasrallah, S. Study of two-dimensional heat and mass transfer during absorption in a metal-hydrogen reactor. *Int. J. Hydrogen Energy* **1995**, *20*, 43–52. [CrossRef]
9. Muthukumar, P.; Prakash Maiya, M.; Srinivasa Murthy, S. Experiments on a metal hydride-based hydrogen storage device. *Int. J. Hydrogen Energy* **2005**, *30*, 1569–1581. [CrossRef]
10. MacDonald, B.D.; Rowe, A.M. Impacts of external heat transfer enhancements on metal hydride storage tanks. *Int. J. Hydrogen Energy* **2006**, *31*, 1721–1731. [CrossRef]
11. Nam, J.; Ko, J.; Ju, H. Three-dimensional modeling and simulation of hydrogen absorption in metal hydride hydrogen storage vessels. *Appl. Energy* **2012**, *89*, 164–175. [CrossRef]
12. Afzal, M.; Mane, R.; Sharma, P. Heat transfer techniques in metal hydride hydrogen storage: A review. *Int. J. Hydrogen Energy* **2017**, *42*, 30661–30682. [CrossRef]
13. Sandrock, G. A panoramic overview of hydrogen storage alloys from a gas reaction point of view. *J. Alloys Compd.* **1999**, *293–295*, 877–888. [CrossRef]
14. Joubert, J.-M.; Paul-Boncour, V.; Cuevas, F.; Zhang, J.; Latroche, M. $LaNi_5$ related AB_5 compounds: Structure, properties and applications. *J. Alloys Compd.* **2021**, *862*, 158163. [CrossRef]

15. Lys, A.; Fadonougbo, J.O.; Faisal, M.; Suh, J.-Y.; Lee, Y.-S.; Shim, J.-H.; Park, J.; Cho, Y.W. Enhancing the hydrogen storage properties of A_xB_y intermetallic compounds by partial substitution: A short review. *Hydrogen* **2020**, *1*, 38–63. [CrossRef]
16. Son, V.B.; Shimkus, Y.Y.; Mozhzhukhin, S.A.; Bocharnikov, M.S.; Fokina, E.E.; Tarasov, B.P. Application of intermetallics (La,Ce)Ni_5 in hydrogen energy storage systems. *Russ. J. Appl. Chem.* **2020**, *93*, 1380–1386. [CrossRef]
17. Tarasov, B.P.; Arbuzov, A.A.; Volodin, A.A.; Fursikov, P.V.; Mozhzhuhin, S.A.; Lototskyy, M.V.; Yartys, V.A. Metal hydride—Graphene composites for hydrogen based energy storage. *J. Alloys Compd.* **2022**, *896*, 162881. [CrossRef]
18. Tarasov, B.; Arbuzov, A.; Mozhzhukhin, S.; Volodin, A.; Fursikov, P.; Davids, M.W.; Adeniran, J.; Lototskyy, M. Metal hydride hydrogen storage (compression) units operating at near-atmospheric pressure of the feed H_2. *Inorganics* **2023**, *11*, 290. [CrossRef]
19. Witman, M.; Allendorf, M.; Stavila, V. Database for Machine Learning of Hydrogen Storage Materials Properties, Version 0.0.5. Zenodo, 2024. Available online: https://zenodo.org/records/10680097 (accessed on 16 November 2026).
20. Hattrick-Simpers, J.R.; Choudhary, K.; Corgnale, C. A simple constrained machine learning model for predicting high-pressure-hydrogen-compressor materials. *Mol. Syst. Des. Eng.* **2018**, *3*, 509–517. [CrossRef]
21. Rahnama, A.; Zepon, G.; Sridhar, S. Machine learning based prediction of metal hydrides for hydrogen storage, part I: Prediction of hydrogen weight percent. *Int. J. Hydrogen Energy* **2019**, *44*, 7337–7344. [CrossRef]
22. Witman, M.; Ling, S.; Grant, D.M.; Walker, G.S.; Agarwal, S.; Stavila, V.; Allendorf, M.D. Extracting an empirical intermetallic hydride design principle from limited data via interpretable machine learning. *J. Phys. Chem. Lett.* **2020**, *11*, 40–47. [CrossRef]
23. Witman, M.D.; Ling, S.; Wadge, M.; Bouzidi, A.; Pineda-Romero, N.; Clulow, R.; Ek, G.; Chames, J.M.; Allendorf, E.J.; Agarwal, S.; et al. Towards Pareto optimal high entropy hydrides via data-driven materials discovery. *J. Mater. Chem. A* **2023**, *11*, 15878–15888. [CrossRef]
24. Wilson, N.; Verma, A.; Maharana, P.R.; Sahoo, A.B.; Joshi, K. HyStor: An experimental database of hydrogen storage properties for various metal alloy classes. *Int. J. Hydrogen Energy* **2024**, *90*, 460–469. [CrossRef]
25. Beltrame, G.; Dematteis, E.M.; Stavila, V.; Rizzi, P.; Baricco, M.; Palumbo, M. Application of machine learning and data augmentation algorithms in the discovery of metal hydrides for hydrogen storage. *Metals* **2025**, *15*, 1221. [CrossRef]
26. Lototskyy, M.V. New model of phase equilibria in metal–hydrogen systems: Features and software. *Int. J. Hydrogen Energy* **2016**, *41*, 2739–2761. [CrossRef]
27. Raissi, M.; Perdikaris, P.; Karniadakis, G.E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **2019**, *378*, 686–707. [CrossRef]
28. Karniadakis, G.E.; Kevrekidis, I.G.; Lu, L.; Perdikaris, P.; Wang, S.; Yang, L. Physics-informed machine learning. *Nat. Rev. Phys.* **2021**, *3*, 422–440. [CrossRef]
29. Rackauckas, C.; Ma, Y.; Martensen, J.; Warner, C.; Zubov, K.; Supekar, R.; Skinner, D.; Ramadhan, A.; Edelman, A. Universal differential equations for scientific machine learning. *arXiv* **2020**, arXiv:2001.04385.
30. Neves, A.M.; Großmann, W.; Puszkiel, J.A.; Warfsmann, J.; Hosseini, V.R.; Passing, M.; Carraro, T.; Klassen, T.; Niggemann, O.; Jepsen, J. Kinetic-model identification in metal-hydride reactions using neural network autoencoder surrogate models. *Energy AI* **2025**, *22*, 100659. [CrossRef]
31. Ramadhan, F.; Suwarno, S.; Hijriawan, M. Experimental Dataset on the Absorption–Desorption Dynamics of an AB2 Metal Hydride Hydrogen Storage Tank for a 300 Watt PEM Fuel Cell. Mendeley Data, V1, 2026. Available online: https://data.mendeley.com/datasets/8km9z62tct/1 (accessed on 16 November 2026).
32. Ong, S.P.; Richards, W.D.; Jain, A.; Hautier, G.; Kocher, M.; Cholia, S.; Gunter, D.; Chevrier, V.L.; Persson, K.A.; Ceder, G. Python Materials Genomics (pymatgen): A robust, open-source python library for materials analysis. *Comput. Mater. Sci.* **2013**, *68*, 314–319. [CrossRef]
33. Ward, L.; Dunn, A.; Faghaninia, A.; Zimmermann, N.E.R.; Bajaj, S.; Wang, Q.; Montoya, J.; Chen, J.; Bystrom, K.; Dylla, M.; et al. Matminer: An open source toolkit for materials data mining. *Comput. Mater. Sci.* **2018**, *152*, 60–69. [CrossRef]
34. Ward, L.; Agrawal, A.; Choudhary, A.; Wolverton, C. A general-purpose machine learning framework for predicting properties of inorganic materials. *npj Comput. Mater.* **2016**, *2*, 16028. [CrossRef]
35. Friedman, J.H. Greedy function approximation: A gradient boosting machine. *Ann. Stat.* **2001**, *29*, 1189–1232. [CrossRef]
36. Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; et al. Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.* **2011**, *12*, 2825–2830. Available online: https://dl.acm.org/doi/10.5555/1953048.2078195 (accessed on 16 November 2026).
37. Avrami, M. Kinetics of phase change. I. General theory. *J. Chem. Phys.* **1939**, *7*, 1103–1112. [CrossRef]
38. Kennedy, M.C.; O'Hagan, A. Bayesian calibration of computer models. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **2001**, *63*, 425–464. [CrossRef]
39. Sobol', I.M. Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. *Math. Comput. Simul.* **2001**, *55*, 271–280. [CrossRef]

40. Saltelli, A.; Annoni, P.; Azzini, I.; Campolongo, F.; Ratto, M.; Tarantola, S. Variance based sensitivity analysis of model output. Design and estimator for the total sensitivity index. *Comput. Phys. Commun.* **2010**, *181*, 259–270. [CrossRef]
41. Herman, J.; Usher, W. SALib: An open-source Python library for sensitivity analysis. *J. Open Source Softw.* **2017**, *2*, 97. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.