NRPA-PK: Retrospective Sparse Pharmacokinetic Sampling for Data-Efficient Modelling

Ali Issa
Institute for Global Health
University College London
London, United Kingdom
ali.issa.21@ucl.ac.uk

Frank Kloprogge
Institute for Global Health
University College London
London, United Kingdom
f.kloprogge@ucl.ac.uk

Tristan Cazenave
LAMSADE
Université Paris Dauphine - PSL
Paris, France
tristan.cazenave@lamsade.dauphine.fr

***Abstract*—Richly sampled pharmacokinetic studies often generate dense, redundant concentration-time datasets. We study retrospective, subject-specific compression of dense individual profiles by selecting a small set of informative sampling times that preserves reconstruction accuracy and drug exposure. We formulate sparse sampling as a combinatorial optimization problem over subsets of observed times and solve it with Nested Rollout Policy Adaptation, a Monte Carlo search method. Each candidate design is evaluated by refitting a simple pharmacokinetic model on the selected samples, reconstructing the full profile, and minimizing a loss that combines weighted root mean squared error and relative area-under-the-curve error. We evaluate the method on two real datasets, indomethacin after intravenous bolus dosing and theophylline after a single oral dose, and one simulated multiple-dose study, using budgets of four, six, and eight samples per subject. Relative to Fisher-informed greedy, random, uniform, and heuristic schedules, the proposed method consistently reduces exposure error to the sub-percent or low single-digit range while maintaining comparable profile reconstruction accuracy and improving upper-tail behavior. In a downstream regression task on the simulated study, compressed profiles selected by the method retain enough information for a small neural network to approach the performance obtained with dense profiles. These results show that dense pharmacokinetic profiles can be compressed into lightweight, machine-learning-ready representations without recollecting data.**

***Index Terms*—pharmacokinetics, sparse sampling, Monte Carlo search, pharmacometrics, machine learning**

I. Introduction

Pharmacokinetics (PK) describes how the body absorbs, distributes and eliminates a drug over time, typically through measurements of drug concentration in blood or plasma [1]. PK studies are often built around relatively dense blood sampling schedules in order to characterise these processes with sufficient resolution. In modern settings, population PK analysis and machine learning (ML) models increasingly operate at scale across many subjects, drugs and dose regimens. When dense sampling grids are used for large cohorts, the resulting concentration–time data can be voluminous and highly redundant, making both data management and downstream computation expensive.

A natural way to reduce data volume is to move from dense to *sparse* sampling designs. In this setting, one seeks to select a small number K of sampling times from a larger set of candidate times for each subject. A good sparse design should allow accurate reconstruction of the individual concentration–time profile and preserve clinically relevant exposure measures such as area under the curve (AUC) while using as few samples as possible. If such designs are available, existing and future PK datasets can be stored and processed as informative sparse profiles rather than dense, highly redundant time series, substantially reducing storage and compute when scaling pharmacometric and machine learning analyses to large multi-drug cohorts.

Traditional sparse PK designs are typically derived from Fisher information matrix (FIM) calculations for a prespecified nonlinear mixed-effects model [2], [3] or from simple heuristic grids. These approaches are mainly prospective and population-level: they seek a shared schedule for future subjects, do not directly optimise subject-level reconstruction metrics such as curve RMSE or relative AUC error, and rely on local approximations around nominal parameter values and assumed between-subject variability. Our setting is different. We consider retrospective, subject-specific compression of already observed dense individual profiles: given a completed rich-sampling PK profile, we choose a small subset of that subject's timepoints that preserves the reconstructed curve and exposure as closely as possible. This yields compact, ML-ready representations of legacy dense studies rather than a prospective trial-design rule.

We make the following contributions:

- We introduce **NRPA-PK**, an NRPA-based framework that selects sparse, subject-specific PK sampling schedules from dense profiles by directly optimising a weighted combination of full-profile RMSE and relative AUC error.
- We evaluate NRPA-PK on two real datasets (indomethacin intravenous bolus and theophylline oral single dose) and one simulated multiple-dose study (MAD2) at budgets of $K = 4, 6, 8$ samples per subject, and compare against a Fisher-informed greedy baseline, random sampling and simple time-based heuristics.
- We show that NRPA-based sparse designs can represent dense PK profiles with only 4–8 samples per subject while preserving AUC to within a few percent and maintaining accurate profiles, and that on MAD2 they

This work was supported by UKRI EPSRC, GlaxoSmithKline, and the PSL–UCL–French Embassy in the UK collaboration programme.

produce compressed representations that let a small neural network approach the performance of models trained on dense profiles, clearly outperform `Greedy_FIM`, and remain competitive with `Random`.

II. Methods

A. Data and pharmacokinetic models

1) Datasets and preprocessing: We analysed three pharmacokinetic datasets: two real single-dose studies and one simulated multiple-dose study. All records were harmonised into a common event–observation format with variables for study identifier, subject, time, dose amount, concentration, compartment, route, and body weight.

The first dataset describes plasma theophylline concentrations following a single oral dose in adult subjects. For each individual, we represented the dosing regimen as a single oral depot dose at time $t = 0$ with amount (mg) computed from the reported dose in $\mathrm{mg/kg}$ and body weight, and we treated all post-dose measurements as observations in the central compartment [4]. The second dataset contains plasma indomethacin concentrations after an intravenous bolus (or short infusion). Each subject received a single 50 mg IV bolus at $t = 0$, modelled as an instantaneous input to the central compartment, with subsequent concentrations recorded as observations; body weight was not available in this study [5].

The third dataset is a simulated multiple ascending dose (MAD2) study with repeated oral dosing and rich sampling in the central compartment. We extracted all plasma concentration–time records and the corresponding oral dosing history, converting concentrations from $\mathrm{ng/mL}$ to $\mathrm{mg/L}$. For each subject we focused on the last dosing interval by identifying the final dose and retaining observations from that time onwards; when at least three observations occurred between 0 and 24 h post-dose we restricted to this window, otherwise we kept all post-last-dose observations but capped the profile at the last six samples to avoid extremely long, irregular tails. The data correspond to the publicly available Multiple Ascending Dose Dataset 2 from the Novartis xgx_v1 project; we retained only the simulated plasma concentrations from the PK compartment and treated the resulting profiles as dense reference curves for the sparse sampling experiments [6].

Across all three studies, preprocessing yielded, for each subject, a dosing history $\{(t_j^{\mathrm{dose}}, \mathrm{AMT}_j)\}$ and observation pairs $\{(t_i, y_i)\}$ in consistent units (mg/L); for the oral studies, pre-dose samples with $t \le 0$ were excluded from the candidate observation set. We then computed simple non-compartmental summaries (AUC and an approximate terminal log-linear slope) [7] on each dense profile to obtain rough per-subject estimates of CL, V and, for oral studies, K_a. These served as seeds and loose bounds for a per-subject weighted least-squares fit of the corresponding one-compartment model, yielding a baseline parameter table (CL_0, V_0, K_{a0}) used for initialisation, box constraints and the reference vector θ_0 in the Fisher-like calculations.

Ethics statement. This study used only previously published de-identified public datasets and simulated data. No new experiments involving human participants or animals were conducted by the authors for this work. The experimental procedures involving human subjects underlying the public datasets analysed in this paper were approved by the appropriate ethics committees or institutional review boards, as reported in the original source publications.

2) Pharmacokinetic models: For indomethacin, which was administered as a single intravenous (IV) bolus, we assumed a one-compartment model with linear elimination. Let t_j and AMT_j denote the time and amount of the j-th IV dose. The concentration in the central compartment is

$$C_{\mathrm{IV}}(t; CL, V) = \sum_j \frac{\mathrm{AMT}_j}{V} \exp\!\left(-\frac{CL}{V}(t - t_j)\right) \mathbb{I}\{t \ge t_j\}, \tag{1}$$

where CL is clearance, V is the apparent volume of distribution, and $\mathbb{I}$ is the indicator function.

For the oral datasets (theophylline and MAD2), we used a one-compartment model with first-order absorption and linear elimination. With elimination rate $K_e = CL/V$ and absorption rate constant K_a, the concentration after oral dosing is

$$C_{\mathrm{oral}}(t; CL, V, K_a) = \sum_j \frac{\mathrm{AMT}_j}{V} \frac{K_a}{K_a - K_e} \times \left(e^{-K_e(t-t_j)} - e^{-K_a(t-t_j)}\right) \mathbb{I}\{t \ge t_j\}. \tag{2}$$

When K_a is numerically close to K_e, we fall back to the usual limiting form with a linear prefactor in time in the implementations to avoid numerical instability. In all studies, the area under the concentration–time curve (AUC) over the observed grid was computed using the trapezoidal rule applied to the sorted observation times and concentrations.

B. Objective and subset evaluation

For each subject, let $\{(t_i, y_i)\}_{i=1}^{M}$ denote the full set of post-processing PK observations, consisting of sampling times t_i and measured concentrations y_i with event flag $\texttt{EVID} = 0$. A sparse sampling design of size K corresponds to a subset of indices $S \subset \{1, \ldots, M\}$ with $|S| = K$, i.e., a selection of K sampling times from the original schedule.

1) Refitting on a subset: Given a candidate subset S, we re-estimate the individual PK parameters using only the selected observations $\{(t_i, y_i)\}_{i \in S}$, while keeping the previously extracted dosing history for that subject fixed.

For indomethacin we estimate (CL, V) using the IV model in (1). For theophylline and MAD2 we estimate (CL, V, K_a) using the oral model in (2). In all cases, parameters are obtained by minimising a weighted least-squares criterion

$$\mathcal{L}_{\text{subset}}(\theta; S) = \sum_{i \in S} \left(\frac{y_i - C(t_i; \theta)}{\max\{C(t_i; \theta), \varepsilon\}}\right)^2, \tag{3}$$

where θ denotes the vector of PK parameters, $C(\cdot;\theta)$ is C_{IV} or C_{oral} depending on the study, and $\varepsilon = 0.05$ mg/L provides a floor on the predicted concentration to stabilise the weights.

The optimisation is performed in log-parameter space using L-BFGS-B. For each subject, we initialise from a per-subject baseline estimate (CL_0, V_0, K_{a0}) and impose broad box constraints centred on these values. For example, for most parameters we constrain the search to approximately one fifth to five times the baseline value (with additional hard lower and upper bounds to avoid implausible extremes) and then exponentiate the solution back to the original scale; K_a is constrained to a physiologically plausible range (0.05–5 h^{-1}). For the IV bolus study, where absorption does not apply, the K_a entry is fixed and not estimated.

2) Full-profile reconstruction and loss: After refitting on S, we reconstruct the full concentration–time profile for that subject at all available observation times by evaluating the model with the subset-based parameter estimates:

$$\hat{y}_i = C(t_i; \hat{\theta}_S), \quad i = 1, \ldots, M.$$

From the reconstructed profile we derive two subject-level error metrics.

First, a weighted root mean squared error (RMSE) over the full profile,

$$\mathrm{RMSE}(S) = \sqrt{\frac{1}{M}\sum_{i=1}^{M}\left(\frac{y_i - \hat{y}_i}{\max\{\hat{y}_i, \varepsilon\}}\right)^2}, \tag{4}$$

which emphasises relative rather than absolute errors and avoids over-weighting very low concentrations.

Second, we quantify the error in overall drug exposure by comparing the AUC of the true and reconstructed profiles. Let

$$\mathrm{AUC}_{\text{true}} = \mathrm{trapz}\big(\{t_i, y_i\}_{i=1}^{M}\big), \quad \mathrm{AUC}_{\text{hat}} = \mathrm{trapz}\big(\{t_i, \hat{y}_i\}_{i=1}^{M}\big), \tag{5}$$

where trapz denotes trapezoidal integration over the sorted time grid. We report the absolute relative AUC error (in percent),

$$|\Delta\mathrm{AUC}(S)|\% = 100 \times \frac{|\mathrm{AUC}_{\text{hat}} - \mathrm{AUC}_{\text{true}}|}{\max\{|\mathrm{AUC}_{\text{true}}|, 10^{-9}\}}. \tag{6}$$

For the real datasets, $\mathrm{AUC}_{\text{true}}$ is thus defined with respect to the dense observed profile and serves as an empirical reference rather than a noise-free pharmacokinetic ground truth.

To rank candidate subsets we combine these two measures into a single scalar score

$$\mathrm{Score}(S) = -\left(\alpha\,\mathrm{RMSE}(S) + \beta\,\frac{|\Delta\mathrm{AUC}(S)|}{100}\right), \tag{7}$$

with fixed weights $\alpha = 1.0$ and $\beta = 0.5$. These values were chosen so that, on the datasets considered, the RMSE and relative AUC terms contribute similar magnitudes to the score and were kept fixed across all studies and sampling budgets. More favourable sparse designs (small RMSE and small AUC error) therefore correspond to higher values of $\mathrm{Score}(S)$.

C. NRPA-based sparse sampling (NRPA-PK)

1) State, actions and terminal condition: Sparse sampling for a single subject is cast as a finite-horizon sequential decision problem. The state is the set of indices that have already been selected, or equivalently a partial sequence $(i_1, \ldots, i_\ell)$ of distinct time indices. At each step the available actions are the remaining indices not yet chosen. An episode starts from the empty set and terminates once K indices have been selected, at which point the subset is $S = \{i_1, \ldots, i_K\}$. The terminal reward is $\mathrm{Score}(S)$ defined in (7); intermediate states receive no reward.

2) NRPA-PK algorithm: We adapt NRPA [8], [9] to the sparse PK design setting. NRPA maintains a real-valued score vector $\phi \in \mathbb{R}^M$ indexed by the candidate time indices. At level $L = 0$, a rollout samples a sequence by repeatedly selecting a legal move m from the current state according to a softmax distribution

$$p_\phi(m \mid \text{state}) \propto \exp(\phi_m), \tag{8}$$

until K distinct indices have been chosen or no moves remain. Higher levels reuse the standard NRPA recursion: a level-L call repeatedly invokes level $L-1$ with the current policy, keeps the best-scoring sequence $S^\star$, and updates ϕ by increasing the log-likelihood of $S^\star$ under the rollout distribution while slightly decreasing the scores of alternative moves. Although NRPA operates on ordered sequences, the reward depends only on the underlying set S, so different permutations correspond to the same sparse design. In all experiments we used $L = 2$, which offered a good balance between solution quality and runtime.

From a combinatorial perspective, the design space is large even for moderate numbers of candidate times. For a subject with M observation times and a budget of K samples, there are $\binom{M}{K}$ distinct sparse designs (unordered subsets) and $M!/(M-K)!$ ordered sequences of length K. For example, with $M = 25$ and $K = 8$, exhaustive search would require evaluating $\binom{25}{8} \approx 1.1 \times 10^6$ subsets. In contrast, our fixed level-2 NRPA configuration (120 iterations per level, 8 restarts) evaluates only $8 \times 120^2 \approx 1.15 \times 10^5$ terminal candidates per subject-budget pair. Each terminal evaluation consists of one PK refit on K selected observations plus one reconstruction on the full M-point profile, so runtime scales linearly with the number of subjects and budgets considered; the dominant cost is repeated PK refitting, while different subjects and budgets can be evaluated independently in parallel. In all experiments we used an unbiased policy (no handcrafted rollout bias), a softmax temperature of 1.0, and fixed pseudo-random seeds for reproducibility.

D. Baselines and experimental setup

1) Baseline methods: We compared NRPA-PK against four per-subject baselines for each sampling budget K. `Greedy_FIM` deterministically constructs a sparse design using a simple local Fisher-information-inspired score: starting from an empty set, it repeatedly adds the timepoint with

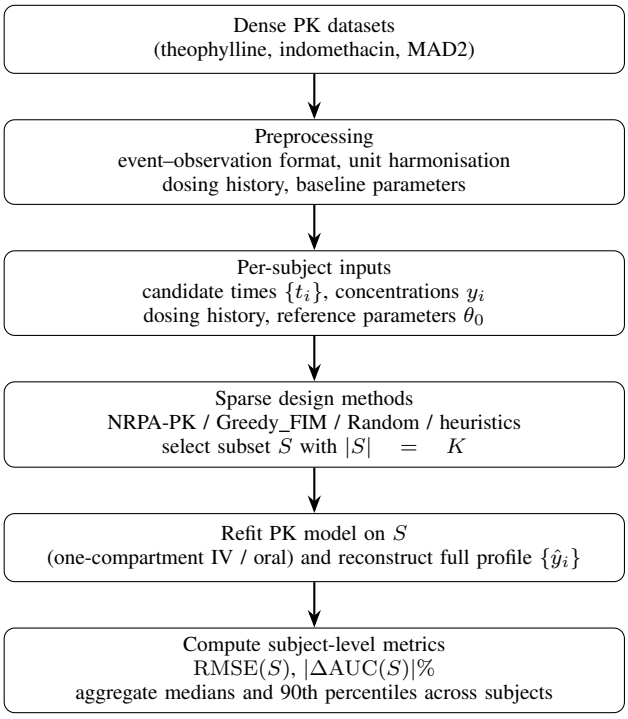

Fig. 1. Overview of the sparse PK sampling design pipeline. NRPA-PK and baselines (Greedy_FIM, Random and simple time-based heuristics) select K-point sparse designs from dense PK profiles, followed by PK refitting and evaluation using RMSE and $|\Delta\mathrm{AUC}|$.

the highest scalar information score (computed from finite-difference sensitivities of concentrations to the PK parameters), penalising times that are very close to already selected ones, until K timepoints have been chosen. `Random` samples K distinct observation times uniformly without replacement, using a fixed pseudo-random seed so that results are reproducible. `UNIFORM` selects approximately equally spaced indices in the sorted time grid for each subject, mimicking an equally spaced sparse sampling grid. `RULE_OF_THUMB` starts from the uniform schedule and forces inclusion of the observed $C_{\max}$ by replacing the nearest non-edge time with the peak concentration time, implementing an "early/peak/late"-style design.

Classical population optimal designs based on the Fisher information matrix optimise an information criterion for population parameters under a mixed-effects model and typically return a single shared sampling grid to be used prospectively for all subjects. In contrast, our setting is retrospective: dense individual profiles and dosing histories are already available, and the goal is to choose subject-specific subsets that best preserve each individual's concentration–time curve and AUC after refitting. Implementing a full population FIM workflow for each dataset would therefore target a different design objective; here `Greedy_FIM` serves as a simple subject-level analogue using local sensitivity information, while `UNIFORM`, `RULE_OF_THUMB` and `Random` represent common time-based or unoptimised sparse grids.

2) Experimental protocol and metrics: We considered sampling budgets $K \in \{4, 6, 8\}$. For each combination of study (`indometh`, `theoph`, `mad2`), subject identifier, method (NRPA-PK, Greedy_FIM, Random, UNIFORM, RULE_OF_THUMB) and K, the pipeline proceeded as follows: for that subject and study, we extracted the relevant dosing history and observation times, selected a subset S of size K according to the method, refit the PK parameters on $\{(t_i, y_i)\}_{i\in S}$, reconstructed the full profile $\{\hat{y}_i\}_{i=1}^{M}$, and finally computed $\mathrm{RMSE}(S)$ and $|\Delta\mathrm{AUC}(S)|\%$.

For each triplet (study, budget K, method) we aggregated results across subjects by reporting the median and 90th percentile of RMSE and relative AUC error. These summary statistics, together with example reconstructed profiles, form the basis of the comparisons in Section III.

E. Downstream ML experiment: AUC regression on MAD2

To test whether NRPA-PK compressed profiles are useful as inputs to downstream machine learning models, we carried out an auxiliary AUC regression experiment on MAD2. For each subject we used the same last-dose 0–24 h window as in the main analysis, defined the target as the trapezoidal area under the concentration–time curve over this interval ($\mathrm{AUC}_{\text{true}}$), and treated prediction of $\mathrm{AUC}_{\text{true}}$ as a supervised regression task.

We compared four input representations at a sampling budget of $K = 4$: (i) `DENSE`, a 25-point regularly sampled concentration vector obtained by interpolating the dense profile onto a fixed 0–24 h grid; (ii) `NRPA-PK`, (iii) `Greedy_FIM` and (iv) `Random`, each represented by a $2K$-dimensional feature vector formed by concatenating the four relative sampling times (scaled to $[0, 1]$ by dividing by 24 h) and the corresponding observed concentrations. For each representation we trained a feed-forward neural network with two hidden ReLU layers (64 and 32 units) and a single linear output, using mean squared error loss [10], standardisation of inputs and targets within each fold, and 5-fold cross-validation across subjects; we report the root mean squared error (RMSE) on $\mathrm{AUC}_{\text{true}}$ and the coefficient of determination R^2 on held-out folds.

III. Results

Table I summarises the reconstruction performance of the main methods (NRPA-PK, Greedy_FIM, Random) across the three studies (indomethacin IV bolus, theophylline oral single dose and the MAD2 multiple-dose simulation) and sampling budgets $K \in \{4, 6, 8\}$. For each combination of study, K and method, we report the median weighted RMSE and the median absolute AUC percentage error ($|\Delta\mathrm{AUC}|$) across subjects; 90th percentile values were also computed to assess upper-tail behaviour but are omitted from the table for compactness. In the table, "NRPA-PK" denotes the NRPA-based method; heuristic UNIFORM and RULE_OF_THUMB baselines are omitted from the table for compactness and their performance is described in Section III-C.

TABLE I
MEDIAN WEIGHTED RMSE AND MEDIAN ABSOLUTE AUC ERROR ($|\Delta\mathrm{AUC}|$, %) ACROSS SUBJECTS. NUMBER OF SUBJECTS: INDOMETHACIN $N = 6$, THEOPHYLLINE $N = 12$, MAD2 $N = 50$.

Study	K	Method	RMSE	$\lvert\Delta\mathrm{AUC}\rvert$
indometh	4	NRPA-PK	0.451	1.3
		Greedy_FIM	2.394	43.3
		Random	0.477	29.7
	6	NRPA-PK	0.451	0.6
		Greedy_FIM	0.853	41.9
		Random	0.428	24.2
	8	NRPA-PK	0.449	2.3
		Greedy_FIM	0.381	35.2
		Random	0.444	18.5
theoph	4	NRPA-PK	0.137	0.2
		Greedy_FIM	0.185	6.6
		Random	0.331	3.9
	6	NRPA-PK	0.136	0.1
		Greedy_FIM	0.139	0.8
		Random	0.156	1.9
	8	NRPA-PK	0.136	0.2
		Greedy_FIM	0.136	1.1
		Random	0.142	1.0
MAD2	4	NRPA-PK	0.024	0.5
		Greedy_FIM	0.026	17.9
		Random	0.028	12.2
	6	NRPA-PK	0.022	0.3
		Greedy_FIM	0.026	11.3
		Random	0.024	11.7
	8	NRPA-PK	0.021	0.7
		Greedy_FIM	0.022	5.9
		Random	0.021	7.0

A. Overall reconstruction performance

For all three studies and sampling budgets, NRPA-PK achieves substantially smaller AUC errors than the baselines while maintaining comparable RMSE. In indomethacin, NRPA-PK keeps median $|\Delta\mathrm{AUC}|$ in the low single digits across $K = 4, 6, 8$, whereas Greedy_FIM and Random often remain in the tens-of-percent range despite similar or larger RMSE. Theophylline shows the same pattern on a smaller absolute scale: NRPA-PK yields sub-percent median AUC errors for all K, with only small RMSE differences between methods at $K = 6, 8$. In MAD2, RMSE is small for all methods, but NRPA-PK again gives the lowest exposure error, with median $|\Delta\mathrm{AUC}|$ below 1% for all budgets versus roughly 6–18% for Greedy_FIM and Random and about 8–15% for the heuristic schedules. Overall, NRPA-PK markedly reduces both median and tail AUC errors while preserving pointwise profile accuracy.

B. Effect of sampling budget

As expected, increasing K generally improves the baselines, especially in AUC error. The largest gains occur when moving from $K = 4$ to $K = 6$, where Greedy_FIM, Random and the heuristic schedules leave the low-budget regime with the largest exposure distortions. For NRPA-PK, however, RMSE changes little with K and median $|\Delta\mathrm{AUC}|$ is already very small at $K = 4$; the best values are typically reached by $K = 6$, with only limited and not strictly monotonic changes at $K = 8$. This indicates that NRPA-PK is most useful in the low-budget regime, where only 4–6 samples per subject are available and naive or myopic designs struggle to control exposure error.

C. Heuristic time-based baselines

The UNIFORM and RULE_OF_THUMB baselines, which approximate simple equally spaced and "early/peak/late" sparse designs, yielded RMSE values similar to the other methods but much larger AUC errors. In the indomethacin IV bolus study, both heuristics produced median $|\Delta\mathrm{AUC}|$ around 70.8% at $K = 4$ and 42–48% at $K = 6, 8$, which is even worse than random sampling (18–30%) and far from NRPA-PK (1.3%, 0.6% and 2.3% at $K = 4, 6, 8$). In the MAD2 multiple-dose simulation, UNIFORM and RULE_OF_THUMB achieved median $|\Delta\mathrm{AUC}|$ between approximately 12–15% at $K = 4$ and 8–11% at $K = 6, 8$, comparable to or slightly better than Greedy_FIM and Random (6–18%) but still an order of magnitude worse than NRPA-PK, which remained below 1% for all budgets.

For the theophylline oral study, the heuristics performed somewhat better but were still clearly dominated by NRPA-PK. At $K = 4$, UNIFORM and RULE_OF_THUMB had median $|\Delta\mathrm{AUC}|$ of about 2.5% and 4.0%, respectively, compared with 0.2% for NRPA-PK; at $K = 6$ and $K = 8$, NRPA-PK kept the median around 0.1–0.2%, while the heuristic designs remained between roughly 1.4–3.5%. Across all three datasets and budgets, these results indicate that simple time-based sparse grids do not reliably control exposure error: although their performance varies by dataset and can sometimes match Greedy_FIM or Random, they remain consistently worse than NRPA-PK and can be dramatically worse in indomethacin.

D. Qualitative reconstructions

Figure 2 illustrates representative reconstructions for one subject from each study at $K = 6$, comparing NRPA-PK, Greedy_FIM and Random designs. For indomethacin, NRPA-PK typically allocates samples in the early post-dose phase and along the declining tail, yielding reconstructed curves that closely follow the dense reference and preserve the elimination slope. Greedy_FIM tends to place more points in regions of high Fisher-like information, sometimes over-emphasising parts of the curve at the expense of late-time exposure, which is consistent with its larger AUC errors. Random designs occasionally miss informative early or late post-dose points, resulting in visibly biased reconstructions and large AUC discrepancies.

For the theophylline and MAD2 oral studies, NRPA-PK tends to select timepoints around absorption peaks and in the post-peak decline (and, for MAD2, around apparent steady-state features), again producing reconstructions that visually match the dense profile while using only 4–8 samples. Because we refit low-dimensional compartmental models by weighted least squares, the reconstructed NRPA-PK curves are not forced to interpolate the selected samples exactly. Some individual measurements, particularly noisy early or terminal points, can therefore lie visibly away from the solid curve even when the overall RMSE and $|\Delta\mathrm{AUC}|$ are small; our objective emphasises global profile shape and exposure rather than exact pointwise fit.

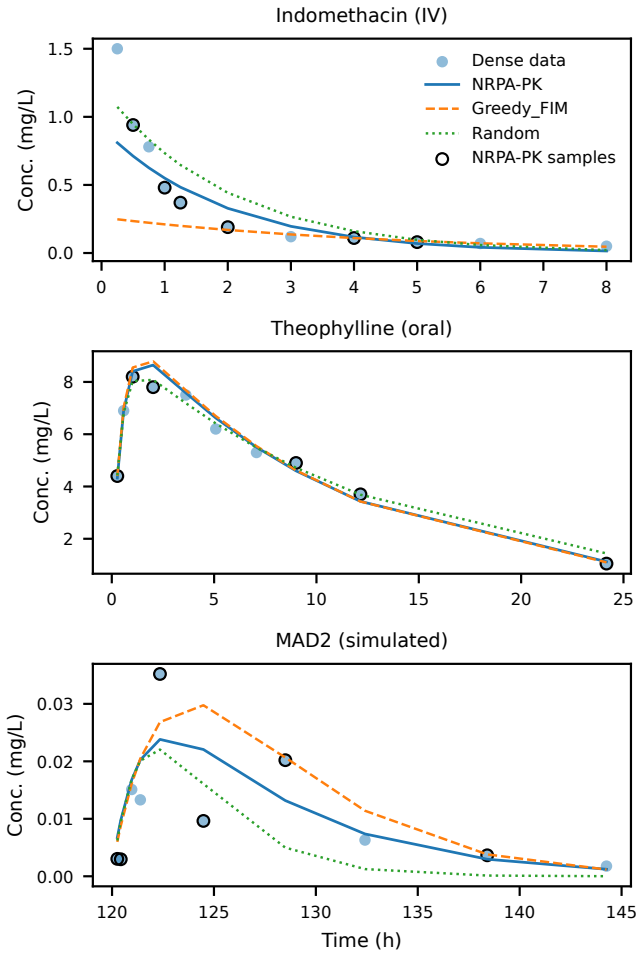

Fig. 2. Representative reconstructions at $K = 6$ for (a) indomethacin, (b) theophylline and (c) MAD2. For each subject, dense observations, NRPA-PK, Greedy_FIM and Random reconstructions are shown; NRPA-PK samples are highlighted.

E. Downstream AUC prediction from compressed MAD2 profiles

Table II summarises the downstream AUC regression performance on the MAD2 dataset at a sampling budget of $K = 4$. As expected, the model trained on dense profiles achieved the lowest AUC RMSE and the highest R^2 (RMSE $\approx$ 0.010, $R^2 \approx 0.99$). Among the sparse representations, NRPA-PK clearly outperformed Greedy_FIM and also improved on Random overall: NRPA-PK gave RMSE ≈ 0.031 and $R^2 \approx 0.89$, compared with RMSE ≈ 0.052, $R^2 \approx 0.62$ for Greedy_FIM and RMSE $\approx$ 0.032, $R^2 \approx 0.80$ for Random. Thus, with only four samples per subject, NRPA-PK retains substantially more information about overall exposure than Greedy_FIM and yields a more predictive sparse representation than an equally sized random schedule.

IV. Discussion

We proposed NRPA-PK, a framework that uses Nested Rollout Policy Adaptation to select sparse, subject-specific PK sampling schedules by directly optimising a reconstruction-based objective. Instead of maximising an analytic information criterion at the population level, NRPA-PK searches over subsets of sampling times for each subject, refits a PK model on those points and evaluates a loss that combines weighted RMSE and relative AUC error. Across two real PK studies and one simulated multiple-dose study, NRPA-PK preserved individual concentration–time profiles and exposure measures with only 4–8 samples per subject and consistently outperformed Fisher-informed, random and heuristic time-based sparse designs in terms of AUC error. We view NRPA-PK as complementary to classical population optimal design methods: it targets retrospective compression of completed dense studies rather than the prospective planning of new sampling schedules.

TABLE II
Downstream AUC regression on MAD2 at $K = 4$: 5-fold cross-validated RMSE and R^2 for a small neural network trained on different subject-level representations.

Representation	RMSE(AUC_{true})	R^2
DENSE	0.0095 $\pm$ 0.0066	0.986 $\pm$ 0.017
NRPA-PK ($K = 4$)	0.0310 $\pm$ 0.0108	0.888 $\pm$ 0.049
Greedy_FIM ($K = 4$)	0.0516 $\pm$ 0.0247	0.620 $\pm$ 0.294
Random ($K = 4$)	0.0323 $\pm$ 0.0151	0.804 $\pm$ 0.161

A. Comparison with baseline designs

The results show a clear pattern across all datasets and sampling budgets. `Greedy_FIM`, which repeatedly adds points with large local Fisher-like information, can achieve reasonable RMSE at higher budgets but often yields large exposure errors, with median $|\Delta\text{AUC}|$ frequently remaining in the 10–40% range and higher upper tails. Random, `UNIFORM` and `RULE_OF_THUMB` schedules behave similarly or worse, especially at low budgets, and their AUC performance is highly variable across subjects. In contrast, NRPA-PK consistently drives median $|\Delta\text{AUC}|$ into the sub-percent to low single-digit range, while keeping RMSE comparable to or better than the baselines and substantially shrinking the upper tail of the AUC error distribution. The large differences in AUC despite similar RMSE in some settings suggest that exposure is more sensitive than pointwise error to deficiencies in sparse designs, particularly when early absorption or terminal elimination phases are under-sampled. By directly optimising a reconstruction-based loss, NRPA-PK avoids designs that fit local segments of the profile at the expense of large global exposure bias.

B. Implications for data-efficient PK and ML

Compressing dense PK profiles into small, information-rich subsets has practical value for both pharmacometrics and machine learning. From a data-engineering perspective, subject-specific sparse designs selected by NRPA-PK can reduce storage and compute requirements when building large multi-drug cohorts, while preserving individual exposure measures that are critical for downstream analyses. In practical terms, dense PK studies already in the archive can be converted into lightweight, ML-ready representations without re-collecting

data, enabling re-use of legacy profiles in larger multi-drug models under fixed storage and compute budgets. From an ML perspective, the MAD2 AUC regression experiment indicates that NRPA-PK subsets retain substantially more information about exposure than equally sized `Greedy_FIM` subsets and modestly more than equally sized random subsets: with only four samples per subject, a neural network trained on NRPA-PK features approaches the performance of a model trained on dense profiles, clearly outperforms `Greedy_FIM`, and is slightly better than `Random` overall, especially in R^2. Although deliberately simple, this experiment suggests that NRPA-PK can serve as a generic compression layer for PK time series in larger pharmacometrics and ML pipelines.

C. Limitations and future work

This study has several limitations. First, we evaluated NRPA-PK on three rich-sampling studies with relatively simple PK models and modest cohort sizes (indomethacin $N = 6$, theophylline $N = 12$, MAD2 $N = 50$). These public benchmarks are standard in PK method development and provide high-resolution profiles that make reconstruction errors easy to interpret, but future work should examine more complex PK–PD structures, dosing schemes and larger archives. The compression problem is essentially per-subject: NRPA-PK is applied independently to each dense profile, so storage and compute savings scale roughly linearly with cohort size; in large multi-drug datasets, the same reduction from dense profiles to $K = 4$–8 samples would yield much larger absolute gains than in these proof-of-concept datasets.

Second, all experiments were retrospective: dense individual profiles were assumed to be available and served as the reference for evaluating sparse designs. Our primary aim was to compress completed studies into compact, ML-ready representations rather than to prescribe prospective sampling schemes. Extending NRPA-PK to a simulation-based prospective design framework, in which candidate sparse schedules are evaluated under a population PK model with variability and model uncertainty, is an important direction for future work.

Third, we focused on two scalar reconstruction metrics (full-profile RMSE and $|\Delta\mathrm{AUC}|$). In many applications, additional endpoints such as $C_{\max}$, $T_{\max}$, partial AUCs over clinically relevant windows or target-attainment criteria may also be of interest. Multi-objective variants of NRPA, such as Pareto-NRPA [11], or formulations that optimise a weighted combination of several PK endpoints, could be used to explore trade-offs between these metrics. Finally, NRPA-PK is more computationally intensive than simple greedy or heuristic strategies because it repeatedly refits the PK model. Parallelisation across subjects, warm-starting the optimisation and using surrogate models for the loss are natural ways to improve scalability, open the door to richer structural models, and enable systematic comparison with classical population optimal design in prospective sparse sampling studies.

V. Conclusion

We introduced NRPA-PK, an NRPA-based framework for retrospective compression of dense individual PK profiles into sparse, subject-specific sampling schedules by directly optimising full-profile reconstruction quality. On two real datasets and one simulated multiple-dose study, NRPA-PK preserved individual exposure with only 4–8 samples per subject and consistently outperformed a Fisher-informed greedy baseline, random sampling and simple time-based heuristics in median absolute $|\Delta\mathrm{AUC}|$, while maintaining comparable weighted RMSE.

These results support NRPA-based search as a practical route to compact, ML-ready PK representations for downstream pharmacometric and machine-learning analyses. Future work will extend the approach to richer PK–PD models, additional clinical endpoints and simulation-based prospective design settings.

References

[1] M. Rowland and T. N. Tozer, *Clinical Pharmacokinetics and Pharmacodynamics: Concepts and Applications*, 4th ed. Philadelphia, PA, USA: Lippincott Williams & Wilkins, 2010.

[2] F. Mentré, A. Mallet, and D. Baccar, "Optimal design in random-effects regression models," *Biometrika*, vol. 84, no. 2, pp. 429–442, 1997.

[3] C. Dumont, S. Chabaud, E. Comets, C. Laffont, and F. Mentré, "Pfim 4.0: An extended, general design evaluation and optimisation tool for nonlinear mixed-effects models," *Computer Methods and Programs in Biomedicine*, vol. 118, no. 2, pp. 138–154, 2015.

[4] R Core Team, "Theoph: Pharmacokinetics of theophylline," `MEMSS` R package documentation, 2025, accessed 2025. [Online]. Available: https://search.r-project.org/CRAN/refmans/MEMSS/html/Theoph.html

[5] ——, "Indometh: Pharmacokinetics of indomethacin," `datasets` R package documentation, 2025, accessed 2025. [Online]. Available: https://search.r-project.org/CRAN/refmans/datasets/html/Indometh.html

[6] Novartis, "xgx_v1 pharmacokinetic datasets (including mad2)," xgxr GitHub repository, 2025, accessed 2025. [Online]. Available: https://github.com/Novartis/xgxr

[7] J. Gabrielsson and D. Weiner, *Pharmacokinetic and Pharmacodynamic Data Analysis: Concepts and Applications*, 5th ed. Stockholm, Sweden: Swedish Pharmaceutical Press, 2016.

[8] T. Cazenave, "Nested monte-carlo search," in *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI'09)*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2009, pp. 456–461. [Online]. Available: https://www.ijcai.org/Proceedings/09/Papers/083.pdf

[9] C. D. Rosin, "Nested rollout policy adaptation for monte carlo tree search," in *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI'11)*. AAAI Press, 2011, pp. 649–654. [Online]. Available: https://www.ijcai.org/Proceedings/11/Papers/115.pdf

[10] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016. [Online]. Available: https://www.deeplearningbook.org/

[11] N. Lallouet, T. Cazenave, and C. Enderli, *Pareto-NRPA: A Novel Monte-Carlo Search Algorithm for Multi-Objective Optimization*. IOS Press, Oct. 2025. [Online]. Available: http://dx.doi.org/10.3233/FAIA251394