



Quick Reference

TDMx Meropenem

Version 1 — 13 August 2026

Author: Sebastian Wicha, Lead Developer

Application

TDMx Meropenem

Developer

Prof. Dr. Sebastian G. Wicha Institute of Pharmacy University of Hamburg Germany
email info (at) tdmx.eu

Copyright © 2015-2026 S.G. Wicha. All Rights Reserved.

Citation

S.G. Wicha, M.G. Kees, A. Solms, I.K. Minichmayr, A. Kratzer, C. Kloft. TDMx – a novel web-based open-access support tool for optimising antimicrobial dosing regimens in clinical routine. Int. J. Antimicrob. Agents. 45: 442-444 (2015).

Purpose

The TDMx Meropenem software is intended for healthcare professionals to calculate individualized dosing regimens for the beta-lactam antibiotic meropenem in adult patients, including critically-ill, ECMO and renal-replacement-therapy (RRT) populations. The software aims to address the problem of high inter-patient variability in meropenem exposure achieved with standard dosing, which can result in concentrations that are not within the range needed for microbiological efficacy.

TDMx Meropenem utilizes patient-specific data such as demographics, laboratory values (including renal function markers) and, if available, previous dosing records and measured meropenem plasma concentrations, to determine individual pharmacokinetic parameters using one of five implemented population pharmacokinetic models covering hospitalised, critically-ill, ECMO and RRT populations (see [Select a model](#)). This computational task is designed to assist healthcare professionals in achieving an adequate percentage of the dosing interval with free (unbound) concentration above the minimum inhibitory concentration ($fT > MIC$) — commonly targeted at 40-100%, and set to $\geq 90\%$ by default in TDMx, reflecting an aggressive target appropriate for critically-ill or carbapenem-resistance-risk patients.

TDMx Meropenem is intended for use in patients prescribed meropenem who require dose individualization for enhanced safety and efficacy, including patients on continuous or intermittent renal replacement therapy. It is designed to be used in clinical settings where such individualized dosing can be monitored and managed.

A decorative graphic at the bottom of the page consisting of several overlapping, wavy lines in shades of blue and orange, creating a sense of movement and depth.

Intended users of the software are healthcare professionals (medical doctors, clinical pharmacists) who have expertise or training in therapeutic drug monitoring and pharmacokinetics. The operating environment for TDMx Meropenem should be healthcare facilities that have the capability to measure drug concentrations and monitor patient responses to antibiotic therapy.

TDMx Meropenem should not be used where it is not feasible to obtain the patient-specific data necessary for dose calculations, or where doing so would contravene the clinical judgement of the treating healthcare professional.

Disclaimer

TDMx is an educational tool and has been created for personal use only. The use of any result generated by TDMx is in any case the sole risk and responsibility of the TDMx user. The information provided by TDMx does not replace clinical judgement. Although TDMx has been validated carefully, there is no guarantee for the accuracy of the provided results. TDMx Meropenem has not been registered as a medical device; dosing decisions should not solely rely on TDMx and remain the responsibility of the treating physician.

Requirements

TDMx is fully tested on Google Chrome® and Mozilla Firefox® browsers. No installation is necessary — TDMx runs entirely in the browser.

Performance characteristics

TDMx was developed using reactive programming with a server-client architecture: TDMx reactively visualizes the user input such as dosing history, patient covariates or drug measurements. TDMx does not require any software to be installed locally.


The population pharmacokinetic models implemented in TDMx Meropenem (see [Select a model](#)) were adapted from published literature models spanning hospitalised, critically-ill (including ECMO) and renal-replacement-therapy patient populations.

Launch TDMx from the web

Go to www.TDMx.eu

Navigate to the menu LAUNCH TDMX and select the module for meropenem.

Startup and disclaimer



During startup of the TDMx app, a welcome dialog is shown with the disclaimer above. Tick “I have read and accept the Disclaimer” and click “Start TDMx” to proceed (Figure 1). Loading a previously saved patient is done afterwards, directly in the main interface (see *Patient File*, below) — not as part of this startup step.

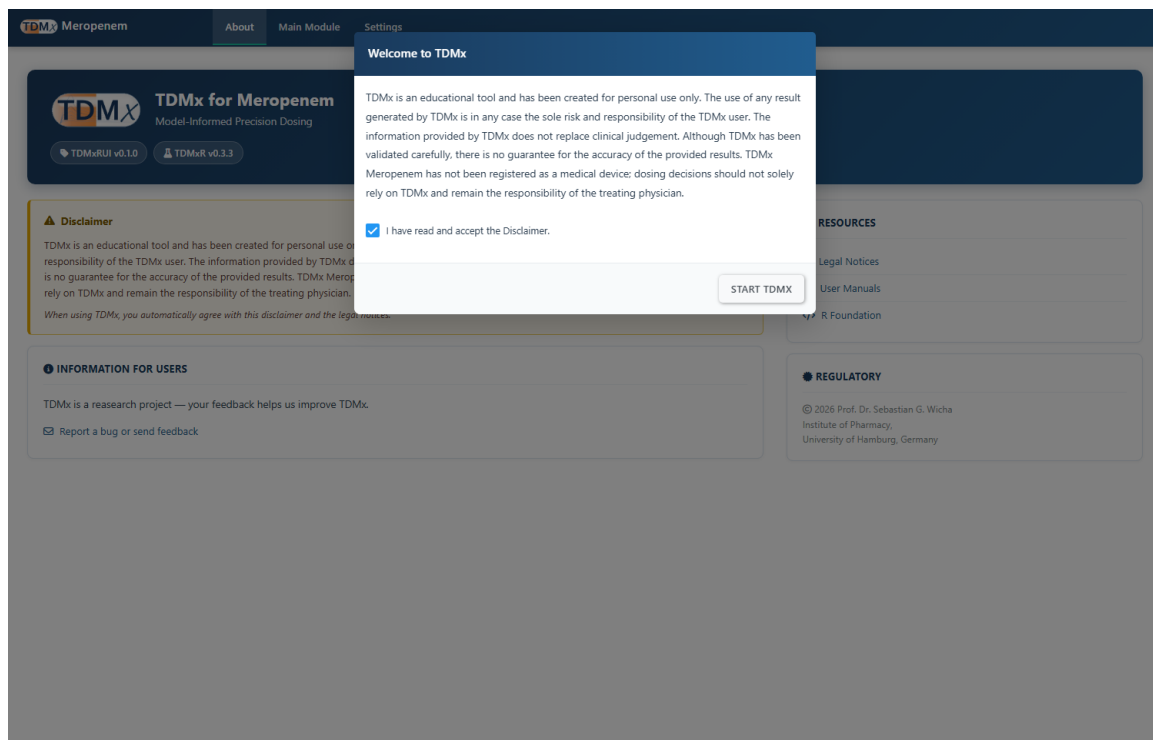


Figure 1: Startup dialog shown when launching TDMx Meropenem.

Main module

The main module consists of two panels: the patient panel (input, left) and the prediction/visualization panel (output, right) (Figure 2).




Figure 2: Main interface of the TDMx web-application, Patient tab, with Individual and Variability ticked.

Input panel (left column): contains three wells — *Patient Record* (patient ID and the Patient/Dosing/Lab/TDM tabs, described below), *Patient File* and *Reference Time*.

- *Patient File*: “LOAD PATIENT” opens a previously saved .tdmx file directly into the current session; “SAVE” downloads the current patient’s full record as a .tdmx file.
- *Reference Time*: the date/time TDMx treats as “now”. Switch to “Use historic time” to backdate a prediction, e.g. when reviewing a case retrospectively.

Output panel (right column): contains graphical and numerical output using the information provided in the input panel, under two tabs — *Single Model* (one selected population PK model, described here) and *Model averaging* (a weighted ensemble of models, see [Model averaging](#)).

In the Single Model tab, tick any combination of the three Options next to the plot:

- Population (PRED, blue): Population-predicted concentration-time data using the dosing record and patient covariates. Observed concentration-time data are not used.

- Individual (IPRED, orange): Individually-predicted concentration-time data using the dosing record, patient covariates and the observed concentration-time data, via Bayesian estimation. Requires at least one TDM measurement.
- Variability: adds shaded prediction-interval bands (up to 95%) around the ticked Population and/or Individual line(s).

PK parameter density plot: whenever “Individual” is ticked and TDM data are available, a small panel of density curves (one per estimated PK parameter, e.g. CL, V1, V2, Q) appears to the right of the main plot. Each curve is the population distribution of that parameter; the shaded band marks the 2.5th–97.5th percentile range and the vertical line marks where this patient’s individually estimated parameter falls, labelled with its percentile. A value near 50% indicates a typical patient; a value near the extremes (close to 2.5% or 97.5%) flags a potential outlier worth a closer look.


Model fit message: also shown once “Individual” is ticked, directly above the numerical table, is a one-line fit message:

- “Model fit acceptable.” (black) — both checks below passed.
- “Poor model fit.” (red) — the relative prediction error between observed and individually-predicted concentrations exceeds the configured threshold (default 40%).
- “...Patient parameters outside population distribution.” (appended, red) — one or more of this patient’s estimated PK parameters falls outside the model’s expected population range (outside its 95% interval by default).

The message turns red if either check fails, so it can also read “Model fit acceptable. Patient parameters outside population distribution.” — treat any red message as a prompt to inspect the fit and parameters carefully rather than accepting the prediction at face value.

Numerical output table:

- Date/Time: Time stamp of the event
- Dose: Administered dose
- Obs.: Observed concentration-time data
- Pop. Pred: Population-predicted concentration-time data using the dosing record and patient covariates. Observed concentration-time data are not used.



- Ind. Pred: Individually-predicted concentration-time data using the dosing record, patient covariates and the observed concentration-time data using Bayesian estimation.
- Creat. Clearance: Creatinine Clearance estimated using the appropriate equation as specified in each implemented population pharmacokinetic model (not computed by the two renal-replacement-therapy models, Shekar 2014 and OJeanson 2021).
- Pop./Ind. CL/V central/V peripheral: Pharmacokinetic parameters of the population or individual model, respectively, considering patient covariates (and, for the individual parameters, the observed concentration-time data). Individual parameters are displayed if “Individual” is selected.
- Pop./Ind. fT>MIC: percentage of the dosing interval with free (unbound) concentration above the minimum inhibitory concentration (MIC), for the population or individual prediction respectively.

Colour coding: the Ind. fT>MIC column is shaded green when the value falls within the target range configured under [Dose optimisation](#) (fT>MIC $\geq 90\%$ by default) and red otherwise — the same colours and threshold used for the “PTA fT>MIC” column in the dose-optimisation tables. This is a quick visual check, not a substitute for clinical judgement — always inspect the actual numbers.

Input panel

The input panel consists of four tabs: Patient, Dosing, Lab and TDM.

Patient

In this tab the user can specify the following (Figure 2):

- A name/identifier for this patient (Patient ID, above the tabs).
- Continuous covariates (Age, Weight, Height) by either typing the value or using the up/down stepper.
- The categorical covariates Sex and Ethnicity, each selected from a dropdown menu.
- A population model for predictions/simulations, or AUTO-SELECT to let TDMx pick the best-fitting model automatically. See [Select a model](#) for details.

Note: not every covariate is used by every model. Hover over a covariate's label to see a tooltip listing which model(s) actually use it for that patient's prediction — a covariate can still be entered even if the currently selected model ignores it.

Covariate	Li 2006	Ehmann 2019	Gijzen 2021	Shekar 2014	OJeanson 2021
Age	used (via Cockcroft-Gault creatinine clearance, plus a direct effect on clearance)	used (via creatinine clearance)	used (via CKD-EPI creatinine clearance)	not used	not used
Weight	used (creatinine clearance & central volume)	used (creatinine clearance & central volume)	used (allometric scaling of clearance and volumes)	not used	not used
Height	not used	not used	not used	not used	not used
Sex	used (via creatinine clearance)	used (via creatinine clearance)	used (via creatinine clearance)	not used	not used
Ethnicity	not used	not used	used (CKD-EPI equation)	not used	not used
Serum creatinine	used (creatinine clearance)	used (creatinine clearance)	used (creatinine clearance)	not used (RRT clears the drug)	not used (RRT clears the drug)

Covariate	Li 2006	Ehmann 2019	Gijssen 2021	Shekar 2014	OJeanson 2021
Albumin	not used	used (central/peripheral volume)	not used	not used	not used
Residual diuresis	n/a	not used	not used	not used	used (clearance)

Dosing

Via the Dosing tab (Figure 3) the user can define all dosing events, and launch dose optimisation.





Figure 3: Input panel to edit dosing-related data.

A new dosing event can be added via “+”, while “-” removes the latest one. Specify the date (dd/mm/yyyy/hh:mm), the dose in mg (or, if switched to “Rate”, the infusion rate in mg/h) and the infusion duration (Dur.) in hours — meropenem is dosed intravenously into the central compartment only, as either a short infusion/bolus or an extended infusion. If a dose is repeated, the number of repeats and the dosing interval (in hours) can be specified via Nr. and Int., respectively.

TDMx offers two dose-optimisation tools; see [Dose optimisation](#) below. For continuous infusion, use the Target Dose Calculation tool’s “continuous” toggle (switched on by default for meropenem), which sets both the dosing interval and the infusion length to 24 h.

Lab

In the Lab tab (Figure 4), the user can specify laboratory data obtained from the patient. These values are used directly as covariates in the population PK model and improve the prediction.






Figure 4: Input panel to enter laboratory data.

Available markers are Serum creatinine, Albumin, Residual diuresis (24 h) and the meropenem minimal inhibitory concentration (MIC). Specify the date of measurement as dd/mm/yyyy/hh:mm and the value in the stated (or an alternative) unit — TDMx automatically converts entered measurements between units. If no measurement is available, leave the field at its default value; if the selected model uses that marker, TDMx substitutes a predefined population value which cancels out its influence. For the MIC, provide an educated guess or assume a worst-case scenario if not (yet) available.

TDM

A new drug plasma concentration can be added via “+”, while “-” removes the latest measurement. Specify the date (dd/mm/yyyy/hh:mm) and the concentration in mg/L

(Figure 5). Untick “Include” to exclude a sample from the Bayesian estimation without removing it from the plot.




Figure 5: Input panel to enter TDM data.

Dose optimisation

Both tools are opened from the Dosing tab and use the selected model, the patient’s covariates and pathogen data (measured or assumed MIC), either with or without drug measurements. Optimising without drug measurements is useful to guide dosing before treatment starts or when no measurements are available; optimising with measurements supports ongoing therapy and is more precise, since individual PK parameters are used.

Note: if “First dose” is selected as the optimisation type, all previous dosing and TDM data will be removed. A warning is shown; accept it only if this is intended.

Dosing Scenarios

Clicking “DOSING SCENARIOS” opens the Dose Optimisation dialog (Figure 6), which evaluates a table of candidate doses (500, 1000 and 2000 mg by default, each compared as a short, near-bolus infusion vs. an extended 4 h infusion, all at an 8 h interval) against a target fT>MIC ($\geq 90\%$ by default). Dose levels, durations and intervals can be edited directly in the table. Click “CALCULATE” to fill in the predicted fT>MIC and probability of target attainment (PTA) for each row, select the most suitable row, and click “ACCEPT DOSE” to add it to the patient’s dosing record.





Figure 6: Dose Optimisation dialog (Dosing Scenarios).

Target Dose Calculation

Clicking “TARGET DOSE CALCULATION” opens the Exact Dose Calculation dialog (Figure 7), which searches for a single dose — within a Min./Max. Dose range (0-6000 mg by default) — that hits one exact target value, either a plasma concentration (Conc.) or a target $ft > MIC$, at a given dosing interval and infusion length. The “continuous” toggle is switched **on by default** for meropenem, setting both the dosing interval and infusion length to 24 h for continuous infusion; untick it to switch to intermittent dosing. Click “CALCULATE” to compute the optimised dose, then “ACCEPT” to add it to the patient’s dosing record.



The screenshot shows the TDMx Meropenem interface with the 'Exact Dose Calculation' dialog box open. The dialog box contains the following sections:

- Timing:** Date/Time of next dose: 13/08/2026/14:00
- Dosing Parameters:**
 - Dosing Interval [h]: 24
 - Infusion length [h]: 24
 - continuous: ☒
 - Add bolus [mg]: 0
- Optimisation Target:**
 - Target Variable: Conc.
 - Target Conc. [mg/L]: 12
 - Min. Dose [mg]: 0
 - Max. Dose [mg]: 6000
- Result:**
 - Optimised Dose [mg]:
 - Optimised Infusion Rate [mg/h]:

Buttons at the bottom of the dialog: CALCULATE, ACCEPT, DISMISS.

The background interface shows the 'Dosing' tab with a table of dosing records and a 'Dose Optimisation' section with 'Add dose' selected. The 'Reference Time' is set to 14/08/2026/11:37.

Figure 7: Exact Dose Calculation dialog (Target Dose Calculation).

Model averaging

Instead of relying on a single population model, the *Model averaging* output tab (Figure 8) combines the predictions of several models into one weighted, individually-estimated prediction — useful when no single model clearly fits best, or to hedge against any one model's misspecification.




Figure 8: Model averaging tab — model-averaged prediction, per-model contributions and model weights.

- **Preset model ensemble:** a dropdown of predefined sets of candidate models — “all models” (all five MER_* models, the default) or “one model” (Li 2006 only, matching the Single Model default).
- **Model ensemble:** the actual list of models included, pre-filled from the preset above but editable — add or remove individual models as needed.
- **Plot:** the thick navy line is the weighted, model-averaged individual prediction; the thin coloured lines are each candidate model’s own individual prediction (same colours as its slice of the “Model Weight” bar on the right). Weights are estimated from how well each model fits this patient’s TDM data.
- **Table:** Ind. Pred./fT>MIC (MAA) columns show the model-averaged individual prediction, colour-coded the same way as the Single Model table (see [Main module](#)).

The weighting criterion (sum of squared residuals by default, or the objective function value) can be changed under [Settings](#). PTA calculation is only available in Single Model mode, not in Model averaging.

Select a model

Model selection (manually)

For all predictions and simulations, a predefined default model — Li 2006 (hospitalised) — is selected. If a different model is more appropriate for a specific patient population, the “Population model” dropdown on the Patient tab lists all five models:

Model	Population	Key covariates used
Li 2006	Hospitalised patients	Weight, Age, Sex, Serum creatinine
Ehmann 2019	Critically-ill patients	Weight, Age, Sex, Serum creatinine, Albumin
Gijssen 2021	Critically-ill patients on ECMO	Weight, Age, Sex, Serum creatinine, Ethnicity
Shekar 2014	Renal replacement therapy, continuous	none of the above (RRT dominates clearance)
OJeanson 2021	Renal replacement therapy, intermittent	Residual diuresis

Model selection (automatically)

TDMx includes an algorithm that automatically selects the best-fitting model for an individual patient. After clicking “AUTO-SELECT”, the algorithm internally compares the predictions of all listed models with the entered TDM data and selects the model with the best fit. The selected model is shown in a pop-up window. For further details on the algorithm, refer to the publication by Uster et al.

Model fit warnings


TDMx continuously checks the fit between the individual prediction and the entered TDM data against two thresholds, and reports the result as the fit message described

under [Main module](#) (“Model fit acceptable.” / “Poor model fit.” / “...Patient parameters outside population distribution.”):

- **Residual check:** flags a poor fit if the relative prediction error, $|\text{predicted} - \text{observed}| / \text{mean}(\text{predicted}, \text{observed})$, exceeds 40% for any TDM sample (configurable).
- **Parameter check:** flags an outlier patient if any individually estimated PK parameter falls outside the 95% interval of that parameter’s population distribution (configurable), i.e. the same percentile shown by the PK parameter density plot.

If either check fails, treat the prediction with extra caution: inspect the plot and PK parameters carefully, and consider that a dosing regimen based on clinical judgement alone might be more appropriate for this patient than the one calculated by TDMx.

Settings



TDMx

Meropenem

AboutMain ModuleSettings

SIMULATION RESOLUTION

HD Resolution [hrs]

SD Resolution [hrs]

0.1

2

Monte Carlo Simulations

Prediction interval method:

Delta-Method

Monte Carlo

Replicates for visualisation

50

250

1,000

Replicates for dose prediction intervals

100

1,000

Variability components in MC simulation:

Inter-individual

Inter-occasion

Residual

Model Averaging

Averaging/selection criterion:

Sum of squared residuals (SSE)

PK TABLE COLOUR CODING

Reference ranges used to colour-code target columns (green/red) in the PK table.

Target fT>MIC [%]

0

90

100

SHARE

Shareable link with patient data encoded in URL:

http://127.0.0.1:3677/?data_json=%7B%22dose_Table%22%3A%5B%7B%22DoseTime%22%3A%2213%2F08%2F2026%22%22%3A%22100%22%22%7D%7D%5D%7D

Figure 9: Settings tab.

To adjust the speed/performance and behaviour of TDMx:

- **Simulation Resolution:** separate time resolutions (in hours) for the high-definition (HD) and standard-definition (SD) plot/table output.
- **Monte Carlo Simulations:** choose the prediction-interval method — Delta-Method (fast, analytical approximation, default) or Monte Carlo (simulation-based) — and, for the latter, the number of replicates used for visualisation and for dose prediction intervals respectively. Variability components (Inter-individual, Inter-occasion, Residual) can be included or excluded from the simulation individually.

- **Model averaging:** the criterion used to weight/select models in Model averaging mode can be changed from the sum of squared residuals (SSE, default) to the objective function value.
- **Share:** generates a shareable link with the current patient's data encoded directly in the URL. This lets a colleague open the same patient record in their own browser session without needing to exchange a *.tdmx file.

Abbreviations

Abbreviation	Meaning
ECMO	Extracorporeal membrane oxygenation
RRT	Renal replacement therapy
Ind.	Individual
HD	High definition
SD	Standard definition (lower time resolution)
PK	Pharmacokinetics
Pred.	Prediction(s)
Obs.	Observed plasma concentration
TDM	Therapeutic drug monitoring
CL	Clearance
V _x	Volume of distribution (1 = central, 2 = peripheral)
CRCL	Creatinine clearance
MIC	Minimum inhibitory concentration

Abbreviation	Meaning
fT>MIC	Percentage of the dosing interval with free (unbound) concentration above the MIC
PTA	Probability of target attainment

