


# Bayesian Workflow for PK/PD modeling using Stan and Torsten

March 2023



**Instructor:** Charles Margossian  
cmargossian@flatironinstitute.org  
<https://charlesm93.github.io/>

## Course description

In this course, we will discuss the tenants of the Bayesian workflow for pharmacokinetic / pharmacodynamic (PK/PD) modeling, and how to execute them using state-of-the-art software. The Bayesian workflow is the iterative process through which we build, fit, and criticize models, with the latter step often motivating useful revisions to our model. We will take advantage of Stan, a Bayesian inference software which boasts a flexible language to specify models, and supports scalable inference algorithms, notably an adaptive Hamiltonian Monte Carlo (HMC) sampler — currently one of the most successful Markov chain Monte Carlo (MCMC) methods. Torsten extends Stan by providing functions to define common PK/PD models and handle the event schedule of clinical trials. This will be a hands-on workshop: students are expected to code and attempt several exercises.

The course is divided into seven modular parts. We will aim to go through as many of them as possible, but our priority is to have an in depth discussion and spend time doing hands-on exercises. You will be given a course package, which includes the course slides, additional documentation, as well as R scripts and Stan files to do the exercises.

### Goals:

- Learn the Stan language, with the Torsten extension
- Understand how to do Bayesian inference with MCMC
- Learn the fundamentals of the Bayesian workflow
- Apply these principles to (population) PK/PD modeling

### Helpful pre-requisites:

- Familiarity with probability (i.e. different types of distributions; conditional distributions; definitions of expectation values, variance, quantiles, etc.)
- Familiarity in a coding language, e.g. R or Python. We will use R scripts for the exercises but the focus will be on coding in Stan. Students are not expected to know Stan.
- Familiarity with PK/PD models (i.e. compartment models; population models; clinical event schedules)

## Installation

To do the exercises and fully take advantage of the workshop, you will need to install Torsten by following the instructions on

<https://metrumresearchgroup.github.io/Torsten/installation/>.

In addition, we will use the R package `cmdstanr` (<https://mc-stan.org/cmdstanr/>), as well as other standard packages such as `ggplot2`. We will send out more detailed instructions on installing Torsten. **I strongly encourage you to install Torsten on your machine a few days before the workshop**, that way, if you encounter an issue you can seek out help, either by contacting me or by using the Stan Discourse (<https://discourse.mc-stan.org/>).

## Course outline

### Morning

1. Fundamentals of Bayesian analysis
  - Specifying a Bayesian model
  - Bayesian inference
  - Beyond inference: the Bayesian modeling workflow

## 2. Algorithms and computational considerations

- Monte Carlo methods
- Hamiltonian Monte Carlo sampling

## 3. Stan

- How it works
- Inference algorithms
- Example 1: Bayesian linear regression
  - *Live demo: writing a model in Stan*
  - Diagnosing the inference algorithm
  - Posterior predictive checks
  - *Exercise 1: improving the model*
  - Computing Markov chains in parallel
  - Variable types
  - Additional language blocks

## 4. Torsten

- Compartment models
- The clinical event schedule
- *Live demo: build, fit, and diagnose a two compartment model*
- *Exercise 3: Parallelize the chains, and do posterior predictive checks*

## Afternoon

## 5. Solving non-linear differential equation (optional)

- Arsenal of tools: matrix exponential, stiff and non-stiff solvers
- *Exercise 4: write, fit, and diagnose a two compartment model using a numerical integrator*
- *Exercise 5: write, fit, and diagnose a nonlinear PK model*
- Solving ODEs in a Bayesian context: challenges and opportunities

## 6. Population models

- Complete pooling, no pooling, and partial pooling
- Principles behind a population model
- *Exercise 6: write, fit, and diagnose a population PK model*

- Divergent transitions and where they come from
- *Exercise 7: reparameterize the model to improve inference*

7. Open discussion and concluding remarks

- Where do Stan and Torsten fit in the pharmacometrician's toolkit?
- What we covered
- What we didn't cover
- Where to learn more