


# Building, fitting, and criticizing Bayesian PK/PD models using Stan and Torsten

April 2022



**Instructor:** Charles Margossian  
charles.margossian@columbia.edu  
<https://charlesm93.github.io/>

The goal of this workshop is to motivate Bayesian modeling, demonstrate its utility across a series of examples, and learn how to use the statistical software Stan and its pharmacometrics extension, Torsten. 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.

**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/>).

## Morning

1. Introduction to 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

## Afternoon

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*
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