

# Probabilistic Programming and Artificial Intelligence

TU Wien, 194.150, VU 6 ECTS

**Jürgen Cito**  
**Markus Böck**

TU Wien, Austria

# Probabilistic Programming and AI: Organization

All information on TISS/TUWEL and website: <https://probprog-ai-tuwien.github.io/2023/>

## Registration

Deadline: October 2nd

Drop-date: October 17th (coincides with deadline for A1)

## Modality/Grading:

6 Lectures, 4 Assignments, 2 Assignment discussions (mandatory), 1 Group project

Grading: 40% Assignments, 60% projects, no exam

## Elective:

066 645 **Data Science**

066 926 **Business Informatics**

066 931 **Logic and Computation**

066 937 **Software Engineering & Internet Computing**

# Probabilistic Programming

1. Represent probability distributions as ~~formulas~~ **programs that generate samples from possible worlds**
2. Build generic algorithms for probabilistic conditioning/inference using probabilistic programs as representations

## Bayesian statistics

- Express statistical assumptions via probability distributions


$$\underbrace{\Pr(\text{parameters, data})}_{\text{joint}} = \underbrace{\Pr(\text{parameters})}_{\text{prior}} \underbrace{\Pr(\text{data} \mid \text{parameters})}_{\text{likelihood}}$$

- Inference from data (via conditioning)

$$\underbrace{\Pr(\text{parameters} \mid \text{data} = \mathbf{x})}_{\text{posterior}}$$

# Probabilistic Programming and AI: What is thinking?

How can we describe the intelligent inferences made in everyday human reasoning?



How can we engineer intelligent machines?

---

## Computational theory of mind



mind = computer




mental representations =  
computer programs

**run(program)**

thinking =  
running a program

# Probabilistic Programming and AI: What kind of program can represent thinking?

Structure



**Knowledge**

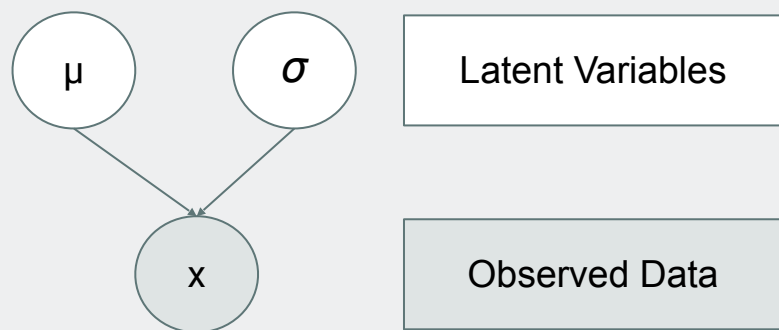
Probability




**Uncertainty**

# Probabilistic Models: Brief Teaser

$$\begin{aligned}\mu &\sim \text{Normal}(\theta, 1\theta) \\ \sigma &\sim \text{LogNormal}(\theta, 5) \\ x &\sim \text{Normal}(\mu, \sigma)\end{aligned}$$





$$\begin{aligned}\mu &\sim \text{Normal}(\theta, 1\theta) \\ \sigma &\sim \text{LogNormal}(\theta, 5) \\ x_1, \dots, x_n &\sim_{iid} \text{Normal}(\mu, \sigma)\end{aligned}$$



# Probabilistic Inference

Run forward

Where will the ball land?

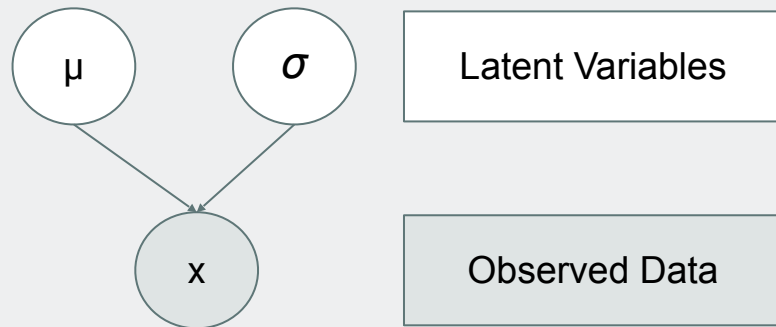


Reason backward

Where did the ball come from?

# Probabilistic Models: Brief Teaser

$$\begin{aligned}\mu &\sim \text{Normal}(\theta, 1\theta) \\ \sigma &\sim \text{LogNormal}(\theta, 5) \\ x &\sim \text{Normal}(\mu, \sigma)\end{aligned}$$



Why are probabilistic models useful?

- Express prior knowledge about the world
- Incorporate noisy data
- Handling uncertainty

Why are Probabilistic Programming Languages useful? **Expressivity!**



# Probabilistic Programming: Expressivity

## Why are PPLs useful? **Expressivity!**

- Probabilistic modeling and inference as first class citizens of a programming language
- Ability to express rich probabilistic models through stochastic control flow (beyond Probabilistic Graphical Models and Bayesian Networks)
- Separating modeling from inference
- **Enable incorporation of programming language and software engineering advances**

# Separating Probabilistic Modeling and Inference

Represent probability distributions by programs that generate samples (simulators)

Using general purpose algorithms for inference using probabilistic programs as model representations

## Probabilistic Model

```
@gen function line_model(xs::Vector{Float64})
    slope = ({:slope} ~ normal(0, 1))
    intercept = ({:intercept} ~ normal(0, 2))

    for (i, x) in enumerate(xs)
        ({:y, i}) ~ normal(slope * x + intercept, 0.1)
    end
    return length(xs)
end;
```

## Inference

```
function do_inference(model, xs, ys, amount_of_computation)
    observations = Gen.choicemap()
    for (i, y) in enumerate(ys)
        observations[{:y, i}] = y
    end

    (trace, _) = Gen.importance_resampling(model, (xs,),
        observations, amount_of_computation);
    return trace
end;

xs = [...];
ys = [...];
sample_trace = Gen.simulate(line_model, (xs, ));
...
trace = do_inference(line_model, xs, ys, 100)
```

# Simple Probabilistic Program: Bayesian Linear Regression in Gen/Julia

```
@gen function line_model(xs::Vector{Float64})
    slope = ({:slope} ~ normal(0, 1))
    intercept = ({:intercept} ~ normal(0, 2))

    for (i, x) in enumerate(xs)
        ({:y, i} ~ normal(slope * x + intercept, 0.1))
    end
    return length(xs)
end;

function do_inference(model, xs, ys, amount_of_computation)
    observations = Gen.choicemap()
    for (i, y) in enumerate(ys)
        observations[{:y, i}] = y
    end

    (trace, _) = Gen.importance_resampling(model, (xs,),
                                           observations, amount_of_computation);
    return trace
end;

xs = [...];
ys = [...];
sample_trace = Gen.simulate(line_model, (xs, ));
...
trace = do_inference(line_model, xs, ys, 100)
```

Slope and Intercept as Latent Variables (Priors)

Likelihood defined as model over slope and intercept with fixed noise

Observed variable samples explicitly recorded in a trace  $\{(:y, i)\}$

Probabilistic Modeling as Generative Function (Simulator)

Observations recalled in inference step and passed to inference procedure

Approximative inference with importance resampling (can be exchanged with MCMC, HMC, NUTS, etc.)

Inference separated from modeling

# Software Engineering for Probabilistic Programming

```
@gen function line_model(xs::Vector{Float64})  
  slope = ({:slope} ~ normal(0, 1))  
  intercept = ({:intercept} ~ normal(0, 2))  
  
  for (i, x) in enumerate(xs)  
    ({:y, i} ~ normal(slope * x + intercept, 0.1))  
  end  
  return length(xs)  
end;  
  
function do_inference(model, xs, ys, amount_of_computation)  
  observations = Gen.choicemap()  
  for (i, y) in enumerate(ys)  
    observations[{:y, i}] = y  
  end  
  
  (trace, _) = Gen.importance_resampling(model, (xs,),  
    observations, amount_of_computation);  
  return trace  
end;  
  
xs = [...];  
ys = [...];  
sample_trace = Gen.simulate(line_model, (xs, ));  
...  
trace = do_inference(line_model, xs, ys, 100)
```

Choice of Priors

Choice of Likelihood

Feedback Cycle

Choice of Inference Method and Parameters

Choice of Domain-specific Visualization

## SE for PPL Research in our research group

Program Comprehension  
(Reasoning about Programs)

Software Evolution  
(Reasoning about Change)

Software Visualization  
(Reasoning about Large-scale Traces)

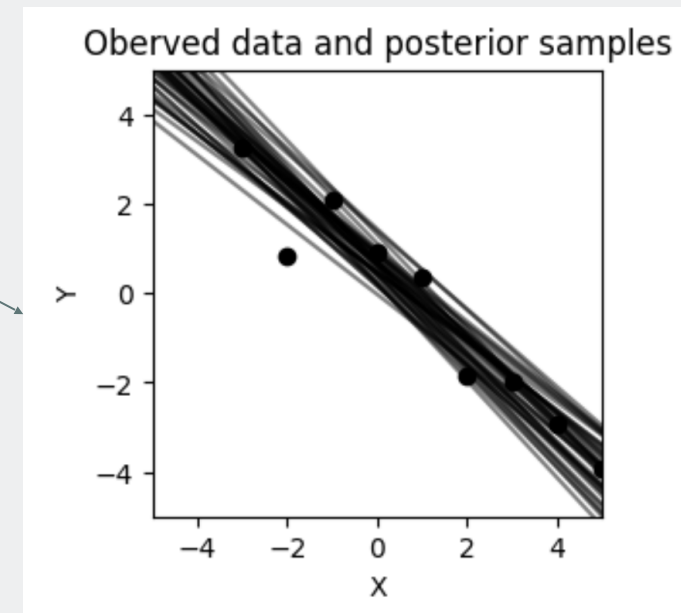
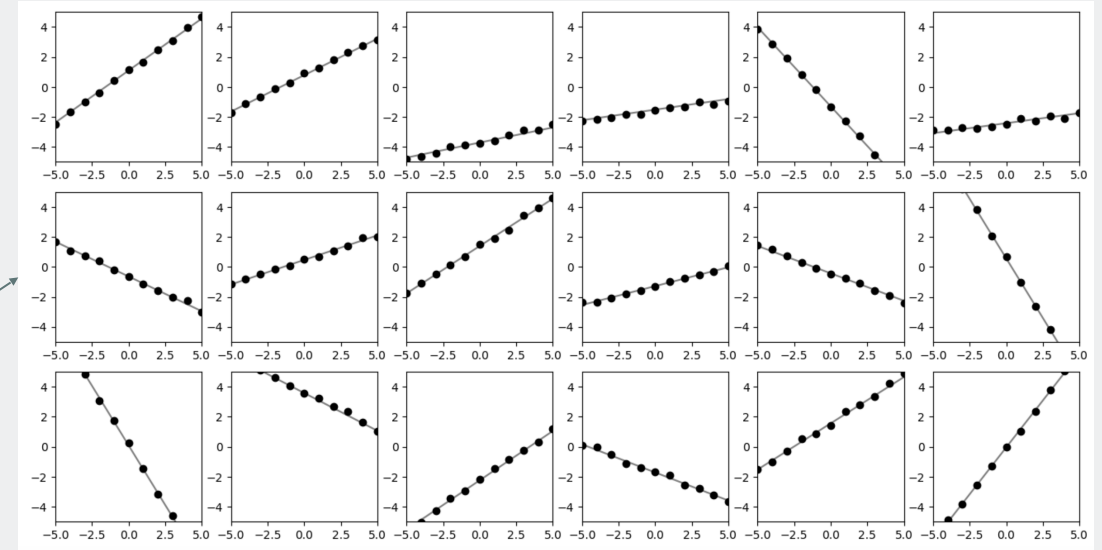
Software Testing  
(Reasoning about Correctness)

# Software Visualization (Reasoning about Traces)

```
function grid(renderer::Function, traces; ncols=6, nrows=3)
    figure(figsize=(16, 8))
    for (i, trace) in enumerate(traces)
        subplot(nrows, ncols, i)
        renderer(trace)
    end
end;
```


```
function overlay(renderer, traces; same_data=true, args...)
    if !isempty(traces)
        renderer(traces[1], show_data=true, args...)
        for i=2:length(traces)
            renderer(traces[i], show_data=!same_data, args...)
        end
    end
end;
```

Support visualizing of generative worlds and posterior distributions that are embedded in domain



# Applications of Probabilistic Programming Languages

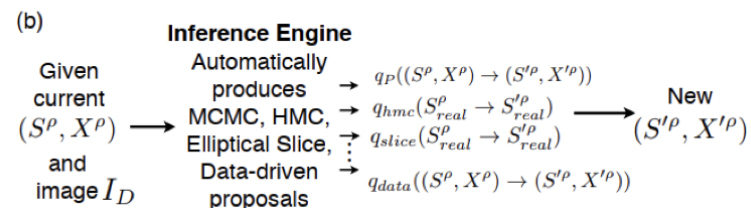
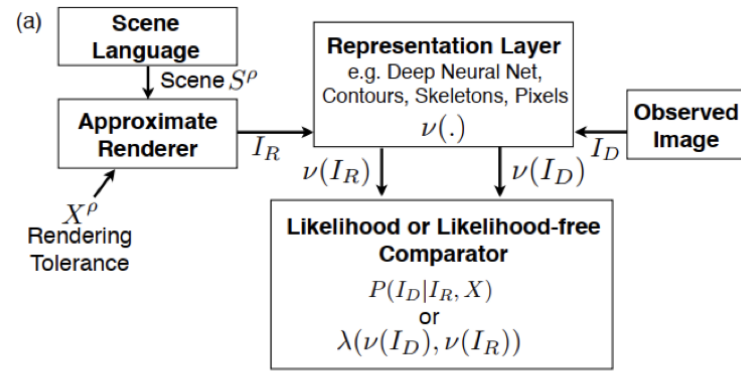
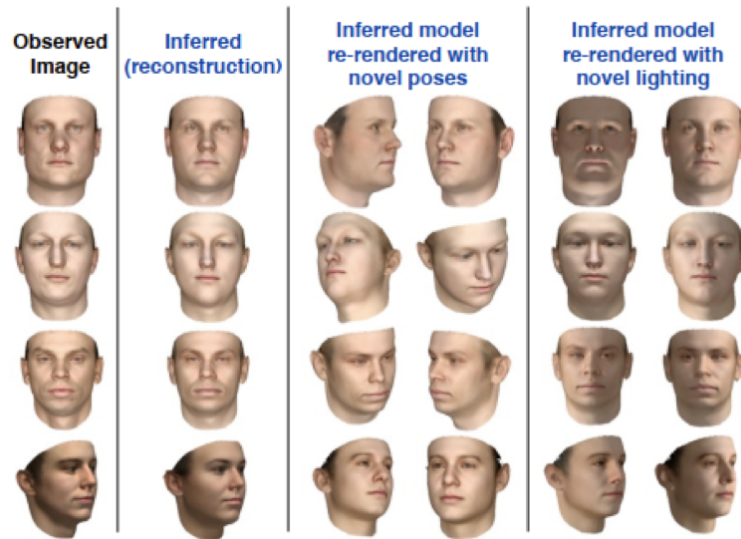
## Inverse Modeling: Extracting 3D structures from images



inference



# Picture: A probabilistic Programming Language for Scene Perception



```

function PROGRAM(MU, PC, EV, VERTEX_ORDER)
# Scene Language: Stochastic Scene Gen
face=Dict();shape = []; texture = [];
for S in ["shape", "texture"]
for p in ["nose", "eyes", "outline", "lips"]
coeff = MvNormal(0,1,1,99)
face[S][p] = MU[S][p]+PC[S][p].*(coeff.*EV[S][p])
end
end
shape=face["shape"][:]; tex=face["texture"][:];
camera = Uniform(-1,1,1,2); light = Uniform(-1,1,1,2)

# Approximate Renderer
rendered_img= MeshRenderer(shape,tex,light,camera)

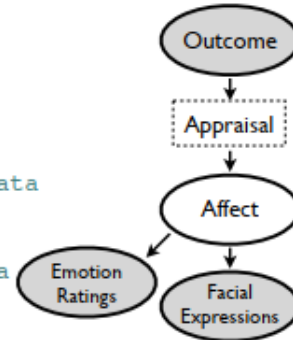
# Representation Layer
ren_ftrs = getFeatures("CNN_Conv6", rendered_img)

# Comparator
#Using Pixel as Summary Statistics
observe(MvNormal(0,0.01), rendered_img-obs_img)
#Using CNN last conv layer as Summary Statistics
observe(MvNormal(0,10), ren_ftrs-obs_cnn)
end

global obs_img = imread("test.png")
global obs_cnn = getFeatures("CNN_Conv6", img)
#Load args from file
TR = trace(PROGRAM, args=[MU, PC, EV, VERTEX_ORDER])
# Data-Driven Learning
learn_datadriven_proposals(TR, 100000, "CNN_Conv6")
load_proposals(TR)
# Inference
infer(TR, CB, 20, ["DATA-DRIVEN"])
infer(TR, CB, 200, ["ELLIPTICAL"])
    
```

# Inferring internal affective states with PPLs

```
1 class MultimodalVAE():
2   def condition(self, outcome, emotion, image):
3     // # generate a new emotion from outcome
4     prediction_mean = self.outcomes_to_affect(outcome)
5     // # sample affect given priors
6     affect = pyro.sample("affect", Normal(prediction_mean, 1))
7     // # generate the facial expression, condition on the observed data
8     face_mean = self.affectToFace_Decoder(affect)
9     face = pyro.sample("face", Bernoulli(face_mean), obs=image)
10    // # generate the outcome ratings, condition on the observed data
11    emo_mean = self.affectToRating_Decoder(affect)
12    emo = pyro.sample("emo", Normal(emo_mean, 1), obs=emotion)
```




Conditioned on:  
low reward




Conditioned on:  
high reward




Conditioned on:  
high anger  
high surprise  
high disappointment



Conditioned on:  
high happiness  
high contentment

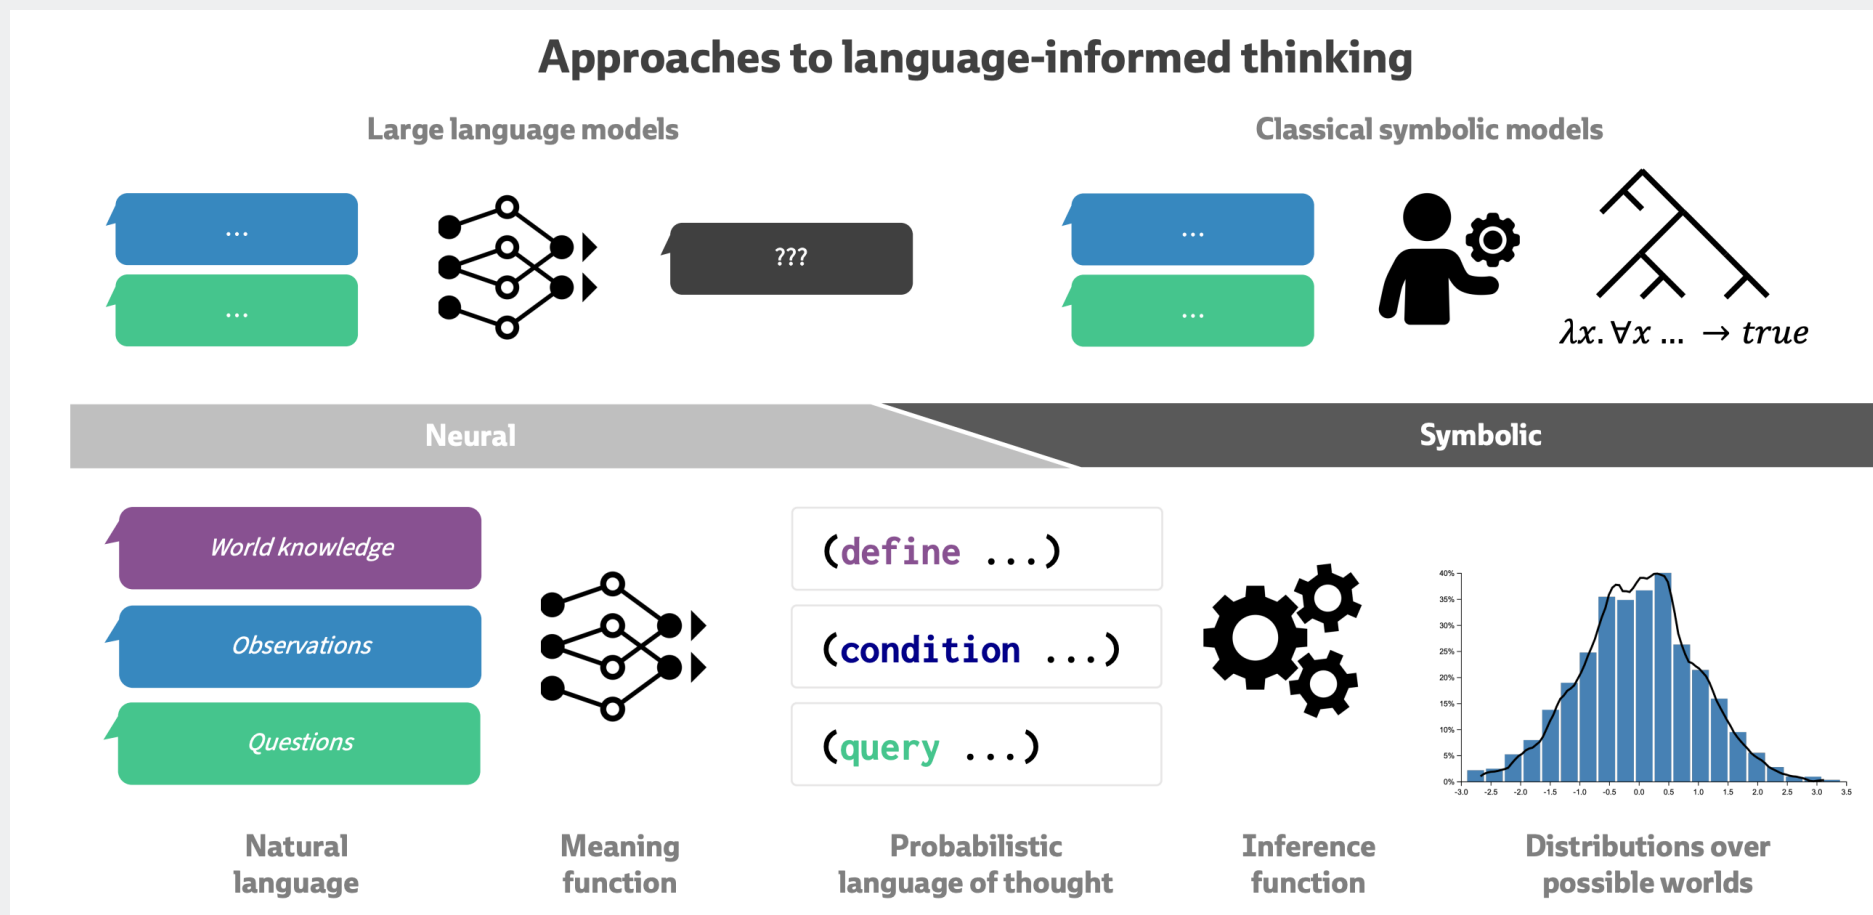


Conditioned on:  
high happiness  
high anger  
high disappointment

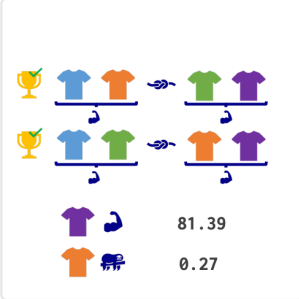
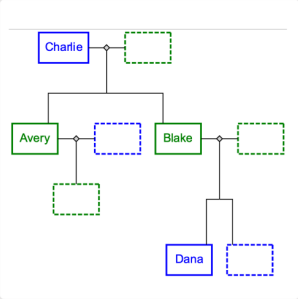

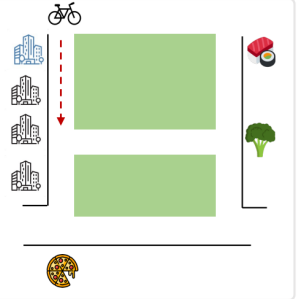





# From Word Models to World Models using PPLs



# From Word Models to World Models using PPLs

	Probabilistic Reasoning	Relational Reasoning	Perceptual and Physical Reasoning	Social Reasoning
	 <p>81.39 0.27</p> <p><i>Bayesian tug-of-war</i></p>	 <p><i>Kinship systems</i></p>	 <p><i>Visual and physical scenes</i></p>	 <p><i>Agents and planning</i></p>
<b>Knowledge about the world</b>	The winner of a match is whichever team is stronger.	A grandfather is the father of one's parent.	Objects vary in both their possible shape and possible color.	Depending on whether they have a bike, people can bike or walk.
<b>Generative world models</b>	<pre>(define won-against (team-1 team-2)   (&gt; (team-strength team-1)       (team-strength team-2)))</pre>	<pre>(define grandfather-of? (name_a name_b)   (exists (lambda (x) (and     (father-of? name_a x)     (parent-of? x name_b))))</pre>	<pre>(define object (obj-id)   (list (choose-shape obj-id)         (choose-color obj-id)))</pre>	<pre>(define actions (agent-id)   (if (has-bike? agent-id)       (list 'is_walking 'is_biking)       (list 'is_walking)))</pre>
<b>Observations about the world</b>	John and Mary faced off against Tom and Sue and won.	Charlie is Dana's grandfather.	There is at least one red mug in this scene.	Alex loves sushi but hates pizza; and he brought his bike to work today.
<b>Condition statements</b>	<pre>(condition (won-against '(john mary)   '(tom sue)))</pre>	<pre>(condition (grandfather-of? 'charlie   'dana))</pre>	<pre>(condition (&gt;= (length   ((filter-shape 'mug)   ((filter-color red)   (objects-in-scene   'this-scene)))) 1))</pre>	<pre>(condition (and (loves? 'alex 'sushi)   (hates? 'alex 'pizza)   (has-bike? 'alex)))</pre>
<b>Questions about the world</b>	Is Mary stronger than Tom?	Which of Charlie's kids is Dana's parent?	How many mugs are there?	What do you think Alex will do?
<b>Query statements</b>	<pre>(query (&gt; (strength 'mary)   (strength 'tom)))</pre>	<pre>(query (filter-tree   (lambda (x) (and     (child-of? x 'charlie)     (parent-of? x 'dana))))))</pre>	<pre>(query (length   ((filter-shape 'mug)   (objects-in-scene   'this-scene))))</pre>	<pre>(query (get_actions 'alex))</pre>

# Probabilistic Programming Languages (PPLs)



# PPL Overview

## Classical PPLs

Have their own modeling Language

Offer a limited set of black-box inference methods

### **Examples:**

Bugs, Stan

# PPL Overview

## Modern PPLs

Are embedded in another language (e.g. Python)

Offer a limited set of black-box inference methods

### Examples:

PyMC, Turing.jl

# PPL Overview

## Deep PPLs

Are embedded in another language (e.g. Python)

Often rely on variational inference and underlying machine learning frameworks (e.g., PyTorch, Tensorflow)

### Examples:

Edward, Pyro

# PPL Overview

## Flexible PPLs

Are embedded in another language (e.g. Python)

Offer the possibility to implement custom inference methods

### **Examples:**

Gen, Beanmachine

# PPL Overview

## Classical PPLs

BUGS  
Stan

## Modern PPLs

PyMC  
Turing.jl

## Deep PPLs

Edward  
Pyro

## Flexible PPLs

Gen  
Beanmachine



# PPL Overview

## Classical PPLs

BUGS

**Stan**

## Modern PPLs

PyMC

**Turing.jl**

## Deep PPLs

Edward

**Pyro**

## Flexible PPLs

**Gen**

**Beanmachine**


# Hello, ~~world~~ coin

Given 10 coinflips (observations): Can we infer the bias of a Coin?

Model:

$\theta \sim \text{Beta}(2,2)$  // coin bias  
 $y \sim \text{Bernoulli}(\theta)$  // coin flip result

Beta(2,2)



# Coin Flip Comparison between different PPLs

Examples can be found at:  
[github.com/ipa-lab/ppl-comparison](https://github.com/ipa-lab/ppl-comparison)

# PPL Comparison

## Mathematical Model:

$\theta \sim \text{Beta}(2,2)$  // coin bias  
 $y \sim \text{Bernoulli}(\theta)$  // coin flip result

# PPL Comparison - Modeling

## Stan

```
stan_model = """
data {
  int N;
  int y[N];
}
parameters {
  real theta;
}
model {
  theta ~ beta(2, 2);
  for (n in 1:N)
    y[n] ~ bernoulli(theta);
}
"""
```

## Gen

```
@gen function my_model(ys::Vector{Bool})
  theta ~ beta(2, 2)
  for (i, y) in enumerate(ys)
    @trace(bernoulli(theta), "y-$i")
  end
end
```

## Turing.jl

```
@model function coinflip(y)
  theta ~ Beta(2, 2)
  N = length(y)
  for n in 1:N
    y[n] ~ Bernoulli(theta)
  end
end
```

## Pyro

```
def simple_model(flips=None):
  a = pyro.param("a", lambda: torch.tensor(2.0))
  b = pyro.param("b", lambda: torch.tensor(2.0))
  theta = pyro.sample("theta", distP.Beta(a,b))

  with pyro.plate("data"):
    return pyro.sample("obs", dist.Bernoulli(theta), obs=flips)
```

## Beanmachine

```
#Heads rate
@bm.random_variable
def theta():
  return dist.Beta(2, 2)

#coin flip
@bm.random_variable
def y(i: int):
  return dist.Bernoulli(theta())
```

# PPL Comparison - Inference

**Turing** (representative for black-box inference)

```
chain = sample(coinflip(data), MH(), 1000)
```

**Pyro** (representative for variational inference)

```
guide = pyro.infer.autoguide.AutoNormal(simple_model)

adam = pyro.optim.Adam({"lr": 0.02}) # Consider decreasing
elbo = pyro.infer.Trace_ELBO()
svi = pyro.infer.SVI(simple_model, guide, adam, elbo)

losses = []
for step in range(1000):
    loss = svi.step(y_obs)
    losses.append(loss)
    if step % 100 == 0:
        print("Elbo loss: {}".format(loss))
```

**Gen** (representative for programmable inference)

```
function my_inference_program(ys::Vector{Bool}, num_iters::Int)
    # Create a set of constraints fixing the
    # y coordinates to the observed y values
    constraints = choicemap()
    for (i, y) in enumerate(ys)
        constraints["y-$i"] = y
    end

    # Run the model, constrained by `constraints`,
    # to get an initial execution trace
    (trace, _) = generate(my_model, (ys,), constraints)


    xs = Float64[]

    # Iteratively update the slope then the intercept,
    # using Gen's metropolis_hastings operator.
    for iter=1:num_iters
        (trace, _) = metropolis_hastings(trace, select(:theta))
        push!(xs, trace[:theta])
    end

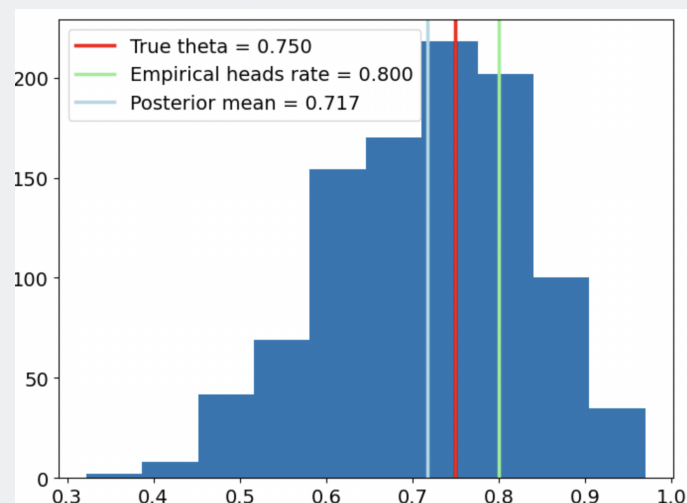
    # From the final trace, read out the slope and
    # the intercept.
    return xs
end
```

# PPL Comparison - Posterior Distributions

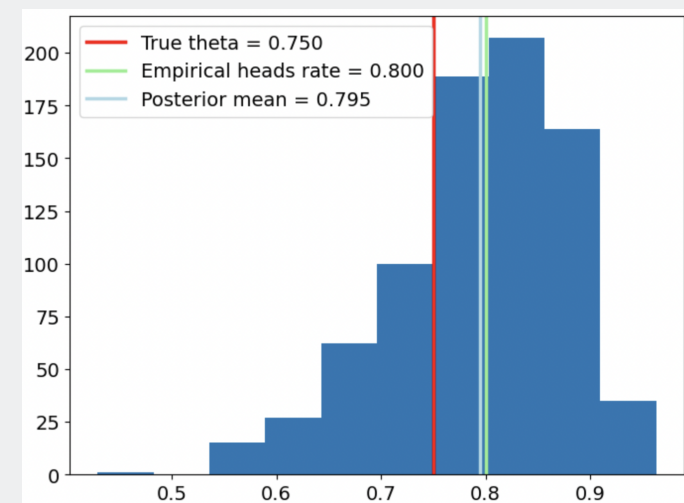
Stan




Beanmachine




Pyro



Turing.jl



Gen



# Probabilistic Programming and AI: Course Overview

Live lectures in Seminarraum FAV 01 A (Seminarraum 183/2) – **Kick-off in FAV Hörsaal 2**  
>> not mandatory, but recommended

04.10. 15:00-17:00	Kick-Off & Lecture 1: <ul style="list-style-type: none"><li>• Introduction to Probabilistic Programming</li><li>• Probability and Bayesian Statistics Primer</li></ul> <b>FAV Hörsaal 2</b>	<ul style="list-style-type: none"><li>• <a href="#">Bayesian Methods for Hackers Probabilistic Programming and Bayesian Inference: Chapter 1</a></li><li>• <a href="#">Seeing Theory: Chapters 1,2,3,5</a></li><li>• <a href="#">3Blue1Brown Bayes Theorem</a></li><li>• <a href="#">AI That Understands the World, Using Probabilistic Programming</a></li></ul>
-----------------------	---	---

04.10. Release Assignment 1 (A1)

11.10. 15:00-17:00	Lecture 2: <ul style="list-style-type: none"><li>• Bayesian Inference and Generative Modelling</li><li>• Probabilistic Programming Languages</li><li>• Implementation Designs</li><li>• Minimal PPL Implementation</li><li>• Independent Sampling</li></ul>	<ul style="list-style-type: none"><li>• <a href="#">Bayesian Inference Framework</a></li><li>• <a href="#">Intuition behind Bayesian inference</a></li><li>• <a href="#">Bayesian posterior sampling</a></li><li>• <a href="#">Generative Models</a></li><li>• <a href="#">A Personal Viewpoint on Probabilistic Programming</a></li></ul>
-----------------------	---	--

11.10. Release Assignment 2 (A2)

18.10. 15:00-17:00	Lecture 3: <ul style="list-style-type: none"><li>• Dependent Sampling</li><li>• Markov Chain Monte Carlo</li><li>• Metropolis Hastings Algorithm</li><li>• Hamiltonian Monte Carlo</li></ul>	<ul style="list-style-type: none"><li>• <a href="#">Why we use dependent sampling to sample from the posterior</a></li><li>• <a href="#">An introduction to the Random Walk Metropolis algorithm</a></li><li>• <a href="#">Paper: Single-Site MH for PPL</a></li><li>• <a href="#">Handbook of MCMC: Chapter 5: MCMC Using Hamiltonian Dynamics</a></li><li>• <a href="#">The intuition behind the Hamiltonian Monte Carlo algorithm</a></li><li>• <a href="#">MCMC Interactive Gallery</a></li><li>• <a href="#">Paper: No-U-Turn Sampler</a></li></ul>
-----------------------	--	--

25.10. 15:00-17:00	Lecture 4: <ul style="list-style-type: none"><li>• Variational Inference</li><li>• Automatic Differentiation VI</li><li>• Stochastic VI</li></ul>	<ul style="list-style-type: none"><li>• <a href="#">KL Divergence - Clearly explained!</a></li><li>• <a href="#">Variational Inference + ELBO Intuition</a></li><li>• <a href="#">Automatic Differentiation and Gradient Descent</a></li><li>• <a href="#">Paper: Automatic Differentiation Variational Inference</a></li><li>• <a href="#">Paper: Stochastic Variational Inference</a></li></ul>
-----------------------	---	---



# Probabilistic Programming and AI: Course Overview

Live lectures in Seminarraum FAV 01 A (Seminarraum 183/2) – **Kick-off in FAV Hörsaal 2**  
 >> not mandatory, but recommended

- 04.10. 15:00-17:00 Kick-Off & Lecture 1:
- Introduction to Probabilistic Programming
  - Probability and Bayesian Statistics Primer

FAV Hörsaal 2

04.10. Release Assignment 1 (A1)

- 11.10. 15:00-17:00 Lecture 2:
- Bayesian Inference and Generative Modelling
  - Probabilistic Programming Languages
  - Implementation Designs
  - Minimal PPL Implementation
  - Independent Sampling

11.10. Release Assignment 2 (A2)

- [Bayesian Methods for Hackers](#)
- [Probabilistic Programming](#)
- [Bayesian Inference](#)
- [Seeing The Forest](#)
- [3Blue1Brown](#)
- [AI That Understands](#)
- [Probabilistic](#)

18.10. 15:00-17:00

Lecture 3:

- Dependent Sampling
- Markov Chain Monte Carlo
- Metropolis Hastings Algorithm

Monte Carlo

- [Why we use dependent sampling to sample from the posterior](#)
- [An introduction to the Random Walk Metropolis algorithm](#)
- [Paper: Single-Site MH for PPL](#)
- [Handbook of MCMC: Chapter 5: MCMC Using Hamiltonian Dynamics](#)
- [The intuition behind the Hamiltonian Monte Carlo algorithm](#)
- [MCMC Interactive Gallery](#)
- [Paper: No-U-Turn Sampler](#)

- 06.12. 15:00-17:00 Lecture 5:
- Advanced Inference
  - TBD

07.12. Office Hours 13:15-15:00 (online)

- 13.12. 15:00-17:00 Lecture 6:
- Probabilistic AI
  - TBD

13.12. A4 Deadline

13.12. Project Proposal Deadline

20.12. 15:00-17:00 Assignment Discussion Session A3 & A4

31.01. 14:00-17:00 Final Project Presentations

Inference  
Differentiation VI

- [KL Divergence - Clearly explained!](#)
- [Variational Inference + ELBO Intuition](#)
- [Automatic Differentiation and Gradient Descent](#)
- [Paper: Automatic Differentiation Variational Inference](#)
- [Paper: Stochastic Variational Inference](#)

# Probabilistic Programming and AI: Course Overview

40% – 4 Assignments with 2 assignment discussion sessions (online and **mandatory!**)

60% – Group project with final presentation

## Assignments:

Assignment 1: Introduction to PPLs

Assignment 2: Minimal PPL implementation

Assignment 3: MH Inference Impl.

Assignment 4: Gradient-based Inference Impl.

**Group project:** You will submit project proposals (we will provide feedback on feasibility)

Initial ideas:

- Applying probabilistic programming to a non-trivial problem
- Implementing an advanced inference algorithm in our minimal PPL
- Implementing a PPL following a different design principle

# Probabilistic Programming and AI: What kind of program can represent thinking?

Structure



Knowledge

Probability



Uncertainty

```
@gen function line_model(xs::Vector{Float64})
  slope = ({:slope} ~ normal(0, 1))
  intercept = ({:intercept} ~ normal(0, 2))

  for (i, x) in enumerate(xs)
    ({:y, i} ~ normal(slope * x + intercept, 0.1))
  end
  return length(xs)
end;

function do_inference(model, xs, ys, amount_of_computation)
  observations = Gen.choicemap()
  for (i, y) in enumerate(ys)
    observations[{:y, i}] = y
  end

  (trace, _) = Gen.importance_resampling(model, (xs,),
    observations, amount_of_computation);
  return trace
end;

xs = [...];
ys = [...];
sample_trace = Gen.simulate(line_model, (xs, ));
...
trace = do_inference(line_model, xs, ys, 100)
```

Slope and Intercept as Latent Variables (Priors)

Likelihood defined as model over slope and intercept with fixed noise

Observed variable samples explicitly recorded in a trace  $\{(:y, i)\}$

Probabilistic Modeling as Generative Function (Simulator)

Observations recalled in inference step and passed to inference procedure

Approximative inference with importance resampling (can be exchanged with MCMC, HMC, NUTS, etc.)

Inference separated from modeling

You can find assignment 1 (A1) on TUWEL

**Remember: Successfully completing A1 until October 17th is mandatory for your final registration**