



Conditional Independence by Typing

MARIA I. GORINOVA, University of Edinburgh

ANDREW D. GORDON, Microsoft Research and University of Edinburgh

CHARLES SUTTON, University of Edinburgh

MATTHIJS VÁKÁR, Utrecht University

A central goal of probabilistic programming languages (PPLs) is to separate modelling from inference. However, this goal is hard to achieve in practice. Users are often forced to re-write their models to improve efficiency of inference or meet restrictions imposed by the PPL. Conditional independence (CI) relationships among parameters are a crucial aspect of probabilistic models that capture a qualitative summary of the specified model and can facilitate more efficient inference.

We present an information flow type system for probabilistic programming that captures conditional independence (CI) relationships and show that, for a well-typed program in our system, the distribution it implements is guaranteed to have certain CI-relationships. Further, by using type inference, we can statically deduce which CI-properties are present in a specified model.

As a practical application, we consider the problem of how to perform inference on models with mixed discrete and continuous parameters. Inference on such models is challenging in many existing PPLs, but can be improved through a workaround, where the discrete parameters are used *implicitly*, at the expense of manual model re-writing. We present a source-to-source semantics-preserving transformation, which uses our CI-type system to automate this workaround by eliminating the discrete parameters from a probabilistic program. The resulting program can be seen as a hybrid inference algorithm on the original program, where continuous parameters can be drawn using efficient gradient-based inference methods, while the discrete parameters are inferred using variable elimination.

We implement our CI-type system and its example application in SlicStan: a compositional variant of Stan.¹

CCS Concepts: • **Theory of computation** → *Random walks and Markov chains*; **Type structures**; *Operational semantics*; • **Mathematics of computing** → **Statistical software**; • **Computing methodologies** → **Learning in probabilistic graphical models**;

Additional Key Words and Phrases: Probabilistic programming, information flow types, static analysis, conditional independence, compiler correctness

¹The implementation is available at <https://github.com/mgorinova/SlicStan>.

Maria Gorinova was supported by the EPSRC Centre for Doctoral Training in Data Science, funded by the UK Engineering and Physical Sciences Research Council (grant EP/L016427/1) and the University of Edinburgh. Matthijs Vákár was funded by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 895827.

Authors' addresses: Maria I. Gorinova, Edinburgh, United Kingdom; email: m.gorinova@ed.ac.uk; A. D. Gordon, Cambridge, United Kingdom; email: adg@microsoft.com; C. Sutton, CA, United States; email: c.sutton@ed.ac.uk; M. Vákár, Utrecht, Netherlands; email: m.i.l.vakar@uu.nl.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2021 Copyright held by the owner/author(s).

0164-0925/2021/12-ART4 \$15.00

<https://doi.org/10.1145/3490421>

ACM Reference format:

Maria I. Gorinova, Andrew D. Gordon, Charles Sutton, and Matthijs Vákár. 2021. Conditional Independence by Typing. *ACM Trans. Program. Lang. Syst.* 44, 1, Article 4 (December 2021), 54 pages. <https://doi.org/10.1145/3490421>

1 INTRODUCTION

The number of **probabilistic programming languages (PPLs)** has grown far and wide, and so has the range of inference techniques they support. Some focus on problems that can be solved analytically and provide a symbolic solution [Gehr et al. 2016]; others are very flexible in the models they can express and use general-purpose inference algorithms [Wood et al. 2014]. Some use gradient-based methods [Carpenter et al. 2017] or message-passing methods [Minka et al. 2014] to provide an efficient solution at the cost of restricting the range of expressible programs. Each option presents its own challenges, whether in terms of speed, accuracy or inference constraints, which is why PPL users often are required to learn a set of model re-writing techniques: to be able to change the program until it can be feasibly used within the backend inference algorithm.

Take for example Stan [Carpenter et al. 2017], which is used by practitioners in a wide range of sciences and industries to analyse their data using Bayesian inference. While efficient inference algorithms exist for continuous-only and for some discrete-only models, it is much less clear what algorithm to use for arbitrary models with large numbers of both discrete and continuous (latent, i.e., unobserved) parameters. Stan has made a conscious choice *not* to support probabilistic models with discrete parameters, to perform inference using (dynamic) **Hamiltonian Monte Carlo (HMC)** [Betancourt and Girolami 2015; Hoffman and Gelman 2014; Neal et al. 2011]), which provides efficient, gradient-based inference for differentiable models. As a result, Stan has often been criticised [Gelman et al. 2015] for its lack of support for discrete parameters. What is usually overlooked is that many models with discrete parameters can, in fact, be accommodated in Stan, by manually marginalising (summing) out the discrete parameters and drawing them conditionally on the continuous parameters [Stan Development Team 2019b, Chapter 7]. One of the core model rewriting techniques is marginalisation: summing over all possible values that a random variable can take to obtain a marginal density function that does not involve that variable. Marginalising efficiently is not always an obvious procedure, as it requires exploiting conditional independence relationships among the variables in the model. For probabilistic graphical models, there are well-known algorithms for enumerating all of the conditional independence assumptions implied by a model. But probabilistic programs are much more general, including control flow and assignment. For this more general case, it is much less clear how to determine conditional independence relationships automatically, and doing so requires combining ideas from traditional program analysis and from probabilistic graphical modelling.

In this article, we introduce an information flow type system that can deduce conditional independence relationships between parameters in a probabilistic program. Finding such relationships can be useful in many scenarios. As an example application, we implement a semantics-preserving source-to-source transformation that automatically marginalises discrete parameters. We work in SlicStan [Gorinova et al. 2019], a form of Stan with a more compositional syntax than the original language. Our system extends SlicStan to support discrete parameters in the case when the discrete parameter space is bounded. This transform corresponds to the *variable elimination* algorithm [Koller and Friedman 2009; Zhang and Poole 1994]: an exact inference algorithm, efficient in models with sparse structure. Combining this transformation with an efficient algorithm for continuous parameters, like HMC, gives us a model-specific, automatically derived inference strategy, which is a composition of variable elimination and the algorithm of choice. While we only focus

on one application in this article, our type system for conditional independence is applicable to program transformations of probabilistic programs more generally, and we believe it can enable other composed-inference strategies.

In short, we make the following contributions:

- (1) *Factorised semantics for SlicStan*: As a basis for proving correctness of our transformation, we extend SlicStan’s type system, so shredding (which slices a SlicStan program into Stan for execution) correctly separates well-typed programs into data preprocessing, main model, and purely generative code (Theorem 1).
- (2) *Main theoretical result*: We show how a very simple, relatively standard information flow type system can be used to capture a conditional independence in probabilistic programs (Section 3) and establish a correspondence between well-typed programs and conditional independence properties of the probability distribution it implements (Theorem 2, Theorem 3).
- (3) *Main practical result*: We describe and implement (in SlicStan) a source-to-source transformation that repeatedly uses the result from (2) to efficiently marginalise out the discrete parameters of the program, and we give a generative procedure for drawing these parameters (Section 4), thus automating inference for mixed discrete-continuous models. We prove that our transformation is semantics-preserving (Theorem 4).

2 SLICSTAN: EXTENDED SYNTAX AND SEMANTICS

SlicStan [Gorinova et al. 2019] is a Stan-like probabilistic programming language. Compared to Stan, it provides extra compositionality by dropping the requirement that programs be block-structured. SlicStan uses type inference in an information-flow type system [Abadi et al. 1999; Gordon et al. 2015; Volpano et al. 1996] to automatically rearrange the program into parts roughly corresponding to the block structure of Stan: pre-processing (data), model, and post-processing (generated quantities). Originally, this *shredding* was developed to compile SlicStan to Stan. In this article, we show that it can be used, more generally, to automatically compile to an efficient program-specific inference scheme.

Like Stan, SlicStan is imperative and allows for deterministic assignment, for-loops, if-statements, probabilistic assignment, and factor-statements. One contribution of this work is that we present an updated version of SlicStan.

A key difference to the original version of SlicStan is the treatment of sampling (\sim) statements. In the original SlicStan paper [Gorinova et al. 2019], a statement such as $x \sim \mathcal{N}(0, 1)$ was understood simply as a syntactic sugar for **factor**($\mathcal{N}(x \mid 0, 1)$): adding a factor to the underlying density of the model, rather than performing actual sampling. In our updated version of SlicStan, sampling statements are part of the core syntax. The semantics of $x \sim \mathcal{N}(0, 1)$ remains *equivalent* to that of **factor**($\mathcal{N}(x \mid 0, 1)$) in terms of density semantics, however it could be *implemented* differently depending on the context. In particular, $x \sim \mathcal{N}(0, 1)$ could be implemented as a simple call to a random number generator in Stan, $x = \mathcal{N}_{rng}(0, 1)$, like in the example in Figure 1.

This way of treating \sim statements differently is useful, as it allows for an increase of the functionality of the SlicStan’s information-flow analysis. Consider, for example, the SlicStan program on the right of Figure 1. Using the original type system, both μ and x_{pred} will be of level **MODEL**, as they are both involved in a \sim statement. Thus, when translated to Stan, both μ and x_{pred} must be inferred with HMC (or another similar algorithm), which is expensive. However, the updated type system of this article allows for x_{pred} to be of level **GENQUANT**, which is preferable: In the context of Stan, this means only μ needs to be inferred with HMC, while x_{pred} can be simply drawn using a random number generator. More generally, the updated SlicStan type system allows for factorising the density defined by the program: For data \mathcal{D} , parameters θ and generated

```

// Stan target program           // SlicStan from POPL'19           // Extended SlicStan
data{ real x; }
parameters{ real  $\mu$ ; }           real  $\mu$ ;
model{ x ~ normal( $\mu$ , 1); }       data real x ~ normal( $\mu$ , 1);
generated quantities{           real x_pred = normal_rng( $\mu$ , 1);
  real x_pred = normal_rng( $\mu$ , 1); }   real x_pred ~ normal( $\mu$ , 1);

```

Fig. 1. Example of difference to previous version of SlicStan.

quantities Q , a program defining a density $p(\mathcal{D}, \theta, Q)$ can be sliced into two programs with densities $p(\mathcal{D}, \theta)$ and $p(Q \mid \mathcal{D}, \theta)$, respectively (Theorem 1). The parameters θ are inferred using HMC (or another general-purpose inference algorithm) according to $p(\mathcal{D}, \theta)$, while the quantities Q are directly generated according to $p(Q \mid \mathcal{D}, \theta)$.

Treating \sim statements differently based on context is very similar in spirit to existing effect-handling-based PPLs [Moore and Gorinova 2018] such as Edward2 and Pyro, where \sim can be handled in different ways. However, in our case, this difference in treatment is determined statically, automatically, and only in the translation to Stan or another backend.

Another difference between Gorinova et al. [2019]’s SlicStan and our updated version is the **target**(S) expression, which we use to capture locally the density defined by statements.

These changes are a small but useful contribution of the current work: They are key to allowing us to decompose the program and compose different inference strategies for efficiency.

In the rest of this section, we give the updated formal syntax, typing, and semantics of SlicStan and describe shredding—the procedure key to the translation of Stan / inference composition.

2.1 Syntax

SlicStan has the following types, programs, L-values, statements, and expressions. We highlight the difference with [Gorinova et al. 2019] with boxes.

SlicStan Types:		SlicStan Program:
$\ell ::= \text{DATA} \mid \text{MODEL} \mid \text{GENQUANT}$	level type	$P ::= \Gamma, S$ program
$n \in \mathbb{N}$	size	
$\tau ::= \text{real} \mid \text{int} \mid \boxed{\text{int}\langle n \rangle} \mid \tau[]$	base type	SlicStan L-Values:
$T ::= (\tau, \ell)$	type	$L ::= x[E_1] \cdots [E_n]$ L-value

SlicStan Typing Environments:

$\Gamma ::= \{x_1 \mapsto T_1, \dots, x_n \mapsto T_n\}$	typing environment
--	--------------------

SlicStan Statements:

SlicStan Statements:	SlicStan Expressions:		
$S ::=$	statement	$E ::=$	expression
$L = E$	assignment	x	variable
$S_1; S_2$	sequence	c	constant
for (x in $E_1 : E_2$) S	for loop	$[E_1, \dots, E_n]$	array
if (E) S_1 else S_2	if statement	$E_1[E_2]$	array element
skip	skip	$f(E_1, \dots, E_n)$	function call
$\boxed{\text{factor}(E)}$	factor statement	$\boxed{[E \mid x \text{ in } E_1 : E_2]}$	array comprehension
$\boxed{L \sim d(E_1, \dots, E_n)}$	sample statement	$\boxed{\text{target}(S)}$	evaluating a density

SlicStan programs consist of a pair Γ, S of a *typing environment* Γ (a finite map that assigns global variables x to their types T) and a *statement* S . Following the usual style of declaring variables in C-like languages, we informally present programs Γ, S in examples by sprinkling the type declarations of Γ throughout the statement S . For example, we write **data real** $x \sim \text{normal}(0, 1)$ for the program $\{x \mapsto (\text{real}, \text{DATA})\}, x \sim \text{normal}(0, 1)$. Sometimes, we will leave out types or write incomplete types in our examples. In this case, we intend for the missing types to be determined using type inference.

As we discuss in detail in Section 2.3, a **factor**(E) statement can be read as multiplying the current *weight* (contribution to the model's joint density) of the program trace by the value of E . Conversely, a **target**(S) expression initialises the weight to 1 and returns the weight that is accumulated after evaluating S . For example, if:

$$\begin{aligned} S &= x \sim \text{normal}(0, 1); y = 2 * x; z \sim \text{normal}(y, 1); \\ &= \text{factor}(\text{normal_pdf}(x|0, 1)); y = 2 * x; \text{factor}(\text{normal_pdf}(z|y, 1)); \end{aligned}$$

Then **target**(S) is semantically equivalent to $\text{normal_pdf}(x|0, 1) * \text{normal_pdf}(z|2 * x, 1)$.

We extend the base types of the language of [Gorinova et al. 2019] with **int**(n), which denotes a positive integer constrained from above by an integer n . For example, if x is of type **int**(2), then x can only be 1 or 2. These types allow us to specify the support of discrete variables, and they can easily be extended to include both upper and lower bounds. For the purpose of our typing rules, we treat **int**(n) identically to **int**. We only differentiate between these types in Section 4, where our transformation uses the size annotation to eliminate a discrete variable.

2.2 Typing

Types T in SlicStan range over pairs (τ, ℓ) of a base type τ , and a level type ℓ . The level types ℓ form a lattice $(\{\text{DATA}, \text{MODEL}, \text{GENQUANT}\}, \leq)$, where $\text{DATA} \leq \text{MODEL} \leq \text{GENQUANT}$. We write $\bigsqcup_{i=1}^n \ell_i$ for the least upper bound of the levels ℓ_1, \dots, ℓ_n . We call variables of level **DATA** *data (variables)*, of level **MODEL** *model parameters*, and of level **GENQUANT** *generated quantities*. We refer to variables that are either of level **MODEL** or **GENQUANT** simply as *parameters*. Given a typing environment Γ , we can consider the well-typedness of expressions and statements, given the types assigned to variables by Γ . The judgment $\Gamma \vdash E : (\tau, \ell)$ means that expression E has type τ and reads only level ℓ and below. The judgment $\Gamma \vdash S : \ell$ means that statement S assigns only to level ℓ and above. We write $\Gamma \vdash S$ as a shorthand for $\Gamma \vdash S : \text{DATA}$.

The typing rules for expressions are those of Gorinova et al. [2019] with added rules for the two constructs of array comprehensions and **target**(S)-expressions. The typing rules for statements are as in Gorinova et al. [2019], with three differences (highlighted in boxes). (**FACTOR**) and (**SAMPLE**) add typing rules for the now new language constructs **factor**(E) and $L \sim d(E_1, \dots, E_n)$. The language supports a finite number of built-in functions f with type $\tau_1, \dots, \tau_n \rightarrow \tau$ and (conditional) distributions $d \in \text{Dist}(\tau_1, \dots, \tau_n; \tau)$ over τ given values of types τ_1, \dots, τ_n .

Typing Rules for Expressions:

(ESUB)	(VAR)	(CONST)	(PRIMCALL)($f : \tau_1, \dots, \tau_n \rightarrow \tau$)
$\Gamma \vdash E : (\tau, \ell) \quad \ell \leq \ell'$	$\Gamma, x : T \vdash x : T$	$\text{ty}(c) = \tau$	$\Gamma \vdash E_i : (\tau_i, \ell_i) \quad \forall i \in 1..n$
$\Gamma \vdash E : (\tau, \ell')$		$\Gamma \vdash c : (\tau, \text{DATA})$	$\Gamma \vdash f(E_1, \dots, E_n) : (\tau, \bigsqcup_{i=1}^n \ell_i)$

$\text{[1.5ex]} \frac{\text{(ARR\text{EL})} \quad \Gamma \vdash E_1 : (\tau[], \ell) \quad \Gamma \vdash E_2 : (\mathbf{int}, \ell)}{\Gamma \vdash E_1[E_2] : (\tau, \ell)}$	$\text{(TARGET)} \quad \frac{\Gamma \vdash S : \ell'' \quad \forall \ell' > \ell. R_{\Gamma+\ell'}(S) = \emptyset^2}{\Gamma \vdash \mathbf{target}(S) : (\mathbf{real}, \ell)}$
$\text{(ARR)} \quad \frac{\Gamma \vdash E_i : (\tau, \ell) \quad \forall i \in 1..n}{\Gamma \vdash [E_1, \dots, E_n] : (\tau[], \ell)}$	$\text{(ARRCOMP)} \quad \frac{\Gamma \vdash E_1 : (\mathbf{int}, \ell) \quad \Gamma \vdash E_2 : (\mathbf{int}, \ell) \quad \Gamma, x : (\mathbf{int}, \ell) \vdash E : (\tau, \ell) \quad x \notin \text{dom}(\Gamma)}{\Gamma \vdash [E \mid x \mathbf{in} E_1 : E_2] : (\tau[], \ell)}$

Typing Rules for Statements:

$\text{(SSUB)} \quad \frac{\Gamma \vdash S : \ell' \quad \ell \leq \ell'}{\Gamma \vdash S : \ell}$	$\text{(ASSIGN)}^3 \quad \frac{\Gamma(L) = (\tau, \ell) \quad \Gamma \vdash E : (\tau, \ell)}{\Gamma \vdash (L = E) : \ell}$	$\text{(IF)} \quad \frac{\Gamma \vdash E : (\mathbf{real}, \ell) \quad \Gamma \vdash S_1 : \ell \quad \Gamma \vdash S_2 : \ell}{\Gamma \vdash \mathbf{if}(E) S_1 \mathbf{else} S_2 : \ell}$
$\text{(SEQ)} \quad \frac{\Gamma \vdash S_1 : \ell \quad \Gamma \vdash S_2 : \ell \quad \mathcal{S}(S_1, S_2) \wedge \mathcal{G}(S_1, S_2)}{\Gamma \vdash (S_1; S_2) : \ell}$	$\text{(FACTOR)} \quad \frac{\Gamma \vdash E : (\mathbf{real}, \mathbf{MODEL})}{\Gamma \vdash \mathbf{factor}(E) : \mathbf{MODEL}}$	$\text{(SKIP)} \quad \frac{}{\Gamma \vdash \mathbf{skip} : \ell}$
$\text{(SAMPLE)}^3 (d \in \text{Dist}(\tau_1, \dots, \tau_n; \tau)) \quad \frac{\Gamma(L) = (\tau, \ell') \quad \Gamma \vdash E_i : (\tau_i, \ell), \forall i \in 1..n \quad \ell = \ell' \sqcup \mathbf{MODEL}}{\Gamma \vdash L \sim d(E_1, \dots, E_n) : \ell}$		
$\text{(FOR)} \quad \frac{\Gamma \vdash E_1 : (\mathbf{int}, \ell) \quad \Gamma \vdash E_2 : (\mathbf{int}, \ell) \quad \Gamma, x : (\mathbf{int}, \ell) \vdash S : \ell \quad x \notin \text{dom}(\Gamma) \quad x \notin W(S)}{\Gamma \vdash \mathbf{for}(x \mathbf{in} E_1 : E_2) S : \ell}$		

In these rules, we make use of the following notation (see Appendix A for precise definitions).

- $W(S)$: the set of variables x that have been assigned to in S .
- $R_{\Gamma+\ell}(S)$: the set of variables x that are read at level ℓ in S .
- $W_{\Gamma+\ell}(S)$: the set of variables x of level ℓ that have been assigned to in S .
- $\widetilde{W}_{\Gamma+\ell}(S)$: the set of variables x of level ℓ that have been \sim -ed in S .
- $W\widetilde{W}_{\Gamma+\ell}(S) = W_{\Gamma+\ell}(S) \cup \widetilde{W}_{\Gamma+\ell}(S)$

The intention in SlicStan is that statements of level ℓ are executed before those of ℓ' if $\ell < \ell'$. To follow that implementation strategy without reordering possibly non-commutative pairs of statements, we impose the condition $\mathcal{S}(S_1, S_2)$ when we sequence S_1 and S_2 in (SEQ).

Definition 1 (Shreddable Seq). $\mathcal{S}(S_1, S_2) \triangleq \forall \ell_1, \ell_2. (\ell_2 < \ell_1) \implies R_{\Gamma+\ell_1}(S_1) \cap W_{\Gamma+\ell_2}(S_2) = \emptyset$.

For example, this excludes the following problematic program:

```
data real sigma = 1;
model real mu ~ normal(0, sigma);
sigma = 2;
```

²We use $\ell' > \ell$ as a shorthand for $\ell \leq \ell' \wedge \neg \ell' \leq \ell$.

³Here, we use $\Gamma(L)$ to look up the type of the L-value L in Γ . Sometimes we will use an overloaded meaning of this notation (Definition 14) to look up the level type of a general expression. Which $\Gamma(\cdot)$ we refer to will be clear from context.

Above, σ and the statements $\sigma=1$ and $\sigma=2$ are of level **DATA**, which means they should be executed before the statement $\mu \sim \text{normal}(\theta, \sigma)$, which is of level **MODEL**. However, this would change the intended semantics of the program, giving μ a $\mathcal{N}(0, 2)$ prior instead of the intended $\mathcal{N}(0, 1)$ prior. This problematic program fails to typecheck in SlicStan, as it is not shreddable: $\neg \mathcal{S}(\mu \sim \text{normal}(\theta, \sigma), \sigma = 2)$.

Definition 2 (Generative Seq). $\mathcal{G}(S_1, S_2) \triangleq \forall \ell \neq \text{MODEL}. \widetilde{W}_{\Gamma-\ell}(S_1) \cap W\widetilde{W}_{\Gamma-\ell}(S_2) = \emptyset \wedge W\widetilde{W}_{\Gamma-\ell}(S_1) \cap \widetilde{W}_{\Gamma-\ell}(S_2) = \emptyset$.

To be able to read $x \sim \mathcal{N}(0, 1)$ at level **GENQUANT**, depending on the context, either as a probabilistic assignment to x or as a density contribution, we impose the condition $\mathcal{G}(S_1, S_2)$ when we sequence S_1 and S_2 . This excludes problematic programs like the following, in which the multiple assignments to y create a discrepancy between the density semantics of the program $p(y) = \mathcal{N}(y \mid 0, 1)\mathcal{N}(y \mid 0, 1)$ and the sampling-based semantics of the program $y = 5$.

```
genquant real y ~ normal(0, 1);
y ~ normal(0, 1);
y = 5;
```

This problematic program fails to typecheck in SlicStan owing to the \mathcal{G} constraint: $\neg \mathcal{G}(y \sim \text{normal}(\theta, 1), y \sim \text{normal}(\theta, 1))$, and also $\neg \mathcal{G}(y \sim \text{normal}(\theta, 1), y = 5)$.

2.3 Operational Semantics of SlicStan Statements

In this article, we use a modified version of the semantics given in Gorinova et al. [2019]. We extend the call-by-value operational semantics given in that paper and derive a more equational form that also includes the generated quantities.

We define a standard big-step operational semantics for SlicStan expressions and statements:

Big-step Relation

$(s, E) \Downarrow V$	expression evaluation
$(s, S) \Downarrow (s', w)$	statement evaluation

Here, s and s' are *states*, V is a *value* and $w \in \mathbb{R}_{>0}$ is a *weight*. Our statements can read and write the state with arbitrary destructive updates. The weight can be thought of as an element of state that stores a positive real value that only gets accessed by multiplying it with the value of an expression E , through the use of **factor**(E)-statements. It can only be read through a **target**(S)-statement that initialises the weight to 1, evaluates the statement S , and returns the final weight.

Formally, states and values are defined as follows:

Values and States:

$V ::=$	value
c	constant
$[V_1, \dots, V_n]$	array
$s ::= x_1 \mapsto V_1, \dots, x_n \mapsto V_n \quad x_i \text{ distinct}$	state (finite map from variables to values)

In the rest of the article, we use the notation for states $s = x_1 \mapsto V_1, \dots, x_n \mapsto V_n$:

- $s[x \mapsto V]$ is the state s , but where the value of x is updated to V if $x \in \text{dom}(s)$, or the element $x \mapsto V$ is added to s if $x \notin \text{dom}(s)$.
- $s[-x]$ is the state s , but where x is removed from the domain of s (if it were present).

We also define lookup and update operations on values:

- If U is an n -dimensional array value for $n \geq 0$ and c_1, \dots, c_n are suitable indexes into U , then the *lookup* $U[c_1] \dots [c_n]$ is the value in U indexed by c_1, \dots, c_n .
- If U is an n -dimensional array value for $n \geq 0$ and c_1, \dots, c_n are suitable indexes into U , then the (functional) *update* $U[c_1] \dots [c_n] := V$ is the array that is the same as U except that the value indexed by c_1, \dots, c_n is V .

The relation \Downarrow is deterministic but partial, as we do not explicitly handle error states. The purpose of the operational semantics is to define a density function in Section 2.4, and any errors lead to the density being undefined. The big-step semantics is defined as follows:

Operational Semantics of Expressions:

$\frac{}{(s, c) \Downarrow c}$	$\frac{V = s(x) \quad x \in \text{dom}(s)}{(s, x) \Downarrow V}$	$\frac{(EVAL ARR) \quad (s, E_i) \Downarrow V_i \quad \forall i \in 1..n}{(s, [E_1, \dots, E_n]) \Downarrow [V_1, \dots, V_n]}$
$\frac{(EVAL ARREL) \quad (s, E_1 \Downarrow V) \quad (s, E_2 \Downarrow c)}{(s, E_1[E_2]) \Downarrow V[c]}$	$\frac{(EVAL PRIMCALL)^4 \quad (s, E_i) \Downarrow V_i \quad \forall i \in 1 \dots n \quad V = f(V_1, \dots, V_n)}{(s, f(E_1, \dots, E_n)) \Downarrow V}$	
$\frac{(EVAL ARRCOMP)^5 \quad (s, E_1) \Downarrow n \quad (s, E_2) \Downarrow m \quad (s, E[i/x]) \Downarrow V_i, \forall n \leq i \leq m}{(s, [E \mid x \text{ in } E_1 : E_2]) \Downarrow [V_n, \dots, V_m]}$		$\frac{(EVAL TARGET) \quad (s, S) \Downarrow (s', w)}{(s, \mathbf{target}(S)) \Downarrow w}$

Operational Semantics of Statements:

$\frac{(EVAL ASSIGN) \text{ (where } L = x[E_1] \dots [E_n]) \quad (s, E_i) \Downarrow V_i \quad \forall i \in 1..n \quad (s, E) \Downarrow V \quad U = s(x) \quad U' = (U[V_1] \dots [V_n] := V)}{(s, L = E) \Downarrow (s[x \mapsto U'], 1)}$		
$\frac{(EVAL SKIP)}{(s, \mathbf{skip}) \Downarrow (s, 1)}$	$\frac{(EVAL SEQ) \quad (s, S_1) \Downarrow (s', w) \quad (s', S_2) \Downarrow (s'', w')}{(s, S_1; S_2) \Downarrow (s'', w * w')}$	$\frac{(EVAL FORFALSE) \quad (s, E_1) \Downarrow c_1 \quad (s, E_2) \Downarrow c_2 \quad c_1 > c_2}{(s, \mathbf{for}(x \text{ in } E_1 : E_2) S) \Downarrow (s, 1)}$
$\frac{(EVAL FORTRUE) \quad \{(s, E_i) \Downarrow c_i\}_{i=1,2} \quad c_1 \leq c_2 \quad (s[x \mapsto c_1], S) \Downarrow (s', w) \quad (s'[-x], \mathbf{for}(x \text{ in } (c_1 + 1) : c_2) S) \Downarrow (s'', w')}{(s, \mathbf{for}(x \text{ in } E_1 : E_2) S) \Downarrow (s'', w * w')}$		
$\frac{(EVAL IFTRUE) \quad (s, E) \Downarrow c \neq 0.0 \quad (s, S_1) \Downarrow (s', w)}{(s, \mathbf{if}(E) S_1 \mathbf{else} S_2) \Downarrow (s', w)}$	$\frac{(EVAL IFFALSE) \quad (s, E) \Downarrow 0.0 \quad (s, S_2) \Downarrow (s', w)}{(s, \mathbf{if}(E) S_1 \mathbf{else} S_2) \Downarrow (s', w)}$	
$\frac{(EVAL FACTOR) \quad (s, E) \Downarrow V}{(s, \mathbf{factor}(E)) \Downarrow (s, V)}$	$\frac{(EVAL SAMPLE)^6 \quad (s, L) \Downarrow V \quad (s, E_i) \Downarrow V_i, \forall 1 \leq i \leq n \quad V' = d(V V_1, \dots, V_n)}{(s, L \sim d(E_1, \dots, E_n)) \Downarrow (s, V')}$	

⁴ $f(V_1, \dots, V_n)$ means applying the built-in function f on the values V_1, \dots, V_n .

⁵Here, we write $E[E'/x]$ for the usual capture avoiding substitution of E' for x in E .

⁶By $d(V|V_1, \dots, V_n)$, we mean the result of evaluating the intended built-in conditional distribution d on V, V_1, \dots, V_n .

Most rules of the big-step operational semantics are standard, with the exception of (**Eval Factor**) and (**Eval Sample**), which correspond to the PPL-specific language constructs **factor** and $L \sim d(E_1, \dots, E_n)$. While we refer to the latter construct as *probabilistic assignment*, its formal semantics is not that of an assignment statement: Both the left- and the right-hand side of the “assignment” are evaluated to a value, for the density contribution $d(V \mid V_1, \dots, V_n)$ to be evaluated and factored into the weight of the current execution trace. Contrary to (**Eval Assign**), there is no binding of a result to a variable in (**Eval Sample**). Of course, as is common in probabilistic programming, it might, at times,⁷ be beneficial to execute these statements as actual probabilistic assignments. Our treatment of these statements is agnostic of such implementation details, however.

The design of the type system ensures that information can flow from a level ℓ to a higher one $\ell' \geq \ell$, but not a lower one $\ell' < \ell$: a noninterference result. To state this formally, we introduce the notions of *conformance between a state s and a typing environment Γ* and *ℓ -equality* of states.

We define a conformance relation on states s and typing environments Γ . A state s *conforms* to an environment Γ , whenever s provides values of the correct types for the variables used in Γ :

Conformance Relation:

$s \models \Gamma$ state s conforms to environment Γ

Rule for the Conformance Relation:

$\frac{\text{(STAN STATE)} \quad V_i \models \tau_i \quad \forall i \in I}{(x_i \mapsto V_i)^{i \in I} \models (x_i : \tau_i)^{i \in I}}$

Here, $V \models \tau$ denotes that the value V is of type τ , and it has the following definition:

- $c \models \mathbf{int}$, if $c \in \mathbb{Z}$, and $c \models \mathbf{real}$, if $c \in \mathbb{R}$.
- $[V_1, \dots, V_n] \models \tau[n]$, if $\forall i \in 1 \dots n. V_i \models \tau$.

Definition 3 (ℓ -equal States). Given a typing environment Γ , states $s_1 \models \Gamma$ and $s_2 \models \Gamma$ are ℓ -equal for level ℓ (written $s_1 \approx_\ell s_2$), if they differ only for variables of a level strictly higher than ℓ :

$$s_1 \approx_\ell s_2 \triangleq \forall x : (\tau, \ell') \in \Gamma. (\ell' \leq \ell \implies s_1(x) = s_2(x)).$$

LEMMA 1 (NONINTERFERENCE OF \vdash). *Suppose $s_1 \models \Gamma$, $s_2 \models \Gamma$, and $s_1 \approx_\ell s_2$ for some ℓ . Then for SlicStan statement S and expression E :*

- (1) *If $\Gamma \vdash E : (\tau, \ell)$ and $(s_1, E) \Downarrow V_1$ and $(s_2, E) \Downarrow V_2$, then $V_1 = V_2$.*
- (2) *If $\Gamma \vdash S : \ell$ and $(s_1, S) \Downarrow s'_1, w_1$ and $(s_2, S) \Downarrow s'_2, w_2$, then $s'_1 \approx_\ell s'_2$.*

PROOF. (1) follows by rule induction on the derivation $\Gamma \vdash E : (\tau, \ell)$, and using that if $\Gamma \vdash E : (\tau, \ell)$, E reads x and $\Gamma(x) = (\tau', \ell')$, then $\ell' \leq \ell$. (2) follows by rule induction on the derivation $\Gamma \vdash S : \ell$ and using (1). We present more details of the proof in Appendix A. \square

2.4 Density Semantics

The semantic aspect of a SlicStan program Γ, S that we are the most interested in is the final weight w obtained after evaluating the program S . This is the value the program computes for the unnormalised joint density $p^*(\mathbf{x}) = p^*(\mathcal{D}, \theta, Q)$ over the data \mathcal{D} , the model parameters θ , and generated

⁷For example, in our Stan backend for SlicStan, if such a statement is of level **MODEL**, it will be executed as density contribution, while if it is of level **GENQUANT**, then it will be executed as a probabilistic assignment.

quantities Q of the program (see Section 2.6). Given a program Γ, S , we separate the typing environment Γ into disjoint parts: Γ_σ and Γ_x , such that Γ_σ contains precisely the variables that are deterministically assigned in S and Γ_x contains those that never get deterministically assigned; that is, the variables \mathbf{x} with respect to which we define the target unnormalised density $p^*(\mathbf{x})$:

$$\Gamma_\sigma = \{(x : T) \in \Gamma \mid x \in W(S)\} \quad \Gamma_x = \Gamma \setminus \Gamma_\sigma.$$

Similarly, any conforming state $s \models \Gamma$ separates as $\sigma \uplus \mathbf{x}$ with

$$\sigma = \{(x \mapsto V) \in s \mid x \in W(S)\} \quad \mathbf{x} = s \setminus \sigma.$$

Then, $\sigma \models \Gamma_\sigma$ and $\mathbf{x} \models \Gamma_x$.

The semantics of a SlicStan program $\Gamma_\sigma, \Gamma_x, S$ is a function $\llbracket S \rrbracket$ on states $\sigma \models \Gamma_\sigma$ and $\mathbf{x} \models \Gamma_x$ that yields a pair of a state σ' and a weight w , such that:

$$\llbracket S \rrbracket(\sigma)(\mathbf{x}) = \sigma', w, \quad \text{where } \sigma \uplus \mathbf{x}, S \Downarrow \sigma' \uplus \mathbf{x}, w.$$

We will sometimes refer only to one of the two elements of the pair σ, w . In those cases, we use the notation: $\llbracket S \rrbracket_s(\sigma)(\mathbf{x})$, $\llbracket S \rrbracket_p(\sigma)(\mathbf{x}) = \llbracket S \rrbracket(\sigma)(\mathbf{x})$. We call $\llbracket S \rrbracket_s$ the *state semantics* and $\llbracket S \rrbracket_p$ the *density semantics* of Γ, S . We will be particularly interested in the density semantics.

The function $\llbracket S \rrbracket_p(\sigma)$ is some positive function $\phi(\mathbf{x})$ of \mathbf{x} . If $\mathbf{x}_1, \mathbf{x}_2$ is a partitioning of \mathbf{x} and $\int \phi(\mathbf{x}) d\mathbf{x}_1$ is finite, we say $\phi(\mathbf{x})$ is an unnormalised density corresponding to the normalised density $p(\mathbf{x}_1 \mid \mathbf{x}_2) = \phi(\mathbf{x}) / \int \phi(\mathbf{x}) d\mathbf{x}_1$ over \mathbf{x}_1 and we write $\llbracket S \rrbracket_p(\sigma)(\mathbf{x}) \propto p(\mathbf{x}_1 \mid \mathbf{x}_2)$. Sometimes, when σ is clear from context, we will leave it implicit and simply write $p(\mathbf{x})$ for $p(\mathbf{x}; \sigma)$.

Next, we observe how the state and density semantics compose.

LEMMA 2 (SEMANTICS COMPOSES). *The state and density semantics compose as follows:*

$$\llbracket S_1; S_2 \rrbracket_s(\sigma)(\mathbf{x}) = \llbracket S_2 \rrbracket_s(\llbracket S_1 \rrbracket_s(\sigma)(\mathbf{x}))(\mathbf{x}) \quad \llbracket S_1; S_2 \rrbracket_p(\sigma)(\mathbf{x}) = \llbracket S_1 \rrbracket_p(\sigma)(\mathbf{x}) \times \llbracket S_2 \rrbracket_p(\llbracket S_1 \rrbracket_s(\sigma)(\mathbf{x}))(\mathbf{x})$$

Throughout the article, we use the following notation to separate the store in a concise way.

Definition 4 ($\Gamma_\ell(s)$ or s_ℓ). For a typing environment Γ and a store $s \models \Gamma$, let $\Gamma_\ell(s) = \{(x \mapsto V) \in s \mid \Gamma(x) = (_, \ell)\}$. When it is clear which typing environment the notation refers to, we write simply s_ℓ instead of $\Gamma_\ell(s)$.

Using this definition, we re-state the noninterference result in the following convenient form:

LEMMA 3 (NONINTERFERENCE OF \vdash REFORMULATED). *Let $\Gamma_\sigma, \Gamma_x \vdash S$ be a well-typed SlicStan program. For all levels $\ell \in \{\text{DATA}, \text{MODEL}, \text{GENQUANT}\}$, there exist unique functions f_ℓ , such that for all $\sigma \models \Gamma_\sigma$, $\mathbf{x} \models \Gamma_x$ and σ' such that $\llbracket S \rrbracket_s(\sigma)(\mathbf{x}) = \sigma'$, $\sigma'_\ell = f_\ell(\{\sigma_{\ell'}, \mathbf{x}_{\ell'} \mid \ell' \leq \ell\})$.*

2.5 Shredding and Translation to Stan

A key aim of SlicStan is to rearrange the input program into three phases of execution, corresponding to the levels of the type system: **DATA** preprocessing, core **MODEL** code to run MCMC or another inference algorithm on, and **GENQUANT**, or generated quantities, which amount to sample post-processing after inference is performed. The motivation for these phases is that they all naturally appear in the workflow of probabilistic programming. The blocks of the Stan are built around this phase distinction, and compilation of SlicStan to Stan and comparable backends requires it.

The phases impose different restrictions on the code and make it incur differing computational costs. The model phase is by far the most expensive to evaluate: Code in this phase tends to be executed repeatedly within the inner loop of an inference algorithm like an MCMC method. Further, it tends to be automatically differentiated [Griewank and Walther 2008] in case gradient-based

inference algorithms are used, which restricts the available programming features and increases the space and time complexity of evaluation. Type inference in SlicStan combined with shredding allows the user to write their code without worrying about the performance of different phases, as code will be shredded into its optimal phase of execution.

The shredding relation is in the core of this rearrangement. Shredding takes a SlicStan statement S and splits it into three *single-level statements* (Definition 5). That is, $S \Downarrow_{\Gamma} S_D, S_M, S_Q$ means we split S into sub-statements S_D, S_M, S_Q , where S_D mentions only **DATA** variables, S_M mentions **DATA** and **MODEL** variables, and S_Q is the rest of the program, and such that the composition $S_D; S_M; S_Q$ behaves the same as the original program S . When combined with type inference, shredding automatically determines optimal statement placement, such that only necessary work is executed in the “heavy-weight” **MODEL** part of inference.

We adapt the shredding from Gorinova et al. [2019], so the following holds for the three sub-statements of a shredded well-typed SlicStan program $\Gamma \vdash S$:

- S_D implements *deterministic data preprocessing*: No contributions to the density are allowed.
- S_M is the *inference core*: It is the least restrictive of the three slices—either or both of S_D and S_Q can be merged into S_M . It can involve contributions to the density that require advanced inference for sampling. Therefore, this is the part of the program that requires the most computation during inference (in Stan, what is run inside HMC).
- S_Q represents *sample post-processing*: Any contributions to the density are generative. That is, they can immediately be implemented using draws from random number generators.

In terms of inference, we can run S_D once as a pre-processing step. Then use a suitable inference algorithm for S_M (in the case of Stan, that is HMC, but we can use other MCMC or VI algorithms), and, finally, we use ancestral sampling for S_Q .⁸

Shredding Relation

$S \Downarrow_{\Gamma} (S_D, S_M, S_Q)$	statement shredding
---	---------------------

Shredding Rules for Statements:

(SHRED ASSIGN)

$$\begin{array}{l} \Gamma(L) = (_, \mathbf{DATA}) \rightarrow S_D = L = E, S_M = S_Q = \mathbf{skip} \\ \Gamma(L) = (_, \mathbf{MODEL}) \rightarrow S_M = L = E, S_D = S_Q = \mathbf{skip} \\ \Gamma(L) = (_, \mathbf{GENQUANT}) \rightarrow S_Q = L = E, S_D = S_M = \mathbf{skip} \\ \hline L = E \Downarrow_{\Gamma} (S_D, S_M, S_Q) \end{array} \quad \begin{array}{l} \text{(SHRED SEQ)} \\ \frac{S_1 \Downarrow_{\Gamma} S_{D_1}, S_{M_1}, S_{Q_1} \quad S_2 \Downarrow_{\Gamma} S_{D_2}, S_{M_2}, S_{Q_2}}{S_1; S_2 \Downarrow_{\Gamma} (S_{D_1}; S_{D_2}), (S_{M_1}; S_{M_2}), (S_{Q_1}; S_{Q_2})} \end{array}$$

(SHRED FACTOR)

$$\begin{array}{l} \Gamma(E) = \mathbf{DATA} \rightarrow S_D = \mathbf{factor}(E), S_M = S_Q = \mathbf{skip} \\ \Gamma(E) = \mathbf{MODEL} \rightarrow S_M = \mathbf{factor}(E), S_D = S_Q = \mathbf{skip} \\ \Gamma(E) = \mathbf{GENQUANT} \rightarrow S_Q = \mathbf{factor}(E), S_D = S_M = \mathbf{skip} \\ \hline \mathbf{factor}(E) \Downarrow_{\Gamma} (S_D, S_M, S_Q) \end{array} \quad \begin{array}{l} \text{(SHRED SKIP)} \\ \frac{}{\mathbf{skip} \Downarrow_{\Gamma} (\mathbf{skip}, \mathbf{skip}, \mathbf{skip})} \end{array}$$

⁸Ancestral (or forward) sampling refers to the method of sampling from a joint distribution by individually sampling variables from the factors constituting the joint distribution. For example, we can sample from $p(x, y) = p(x)p(y | x)$ by randomly generating \hat{x} according to $p(x)$, and then randomly generating \hat{y} according to $p(y | x = \hat{x})$.

(SHRED SAMPLE)

$$\frac{\begin{array}{l} \Gamma(L, E_1, \dots, E_n) = \text{DATA} \rightarrow S_D = L \sim d(E_1, \dots, E_n), S_M = S_Q = \text{skip} \\ \Gamma(L, E_1, \dots, E_n) = \text{MODEL} \rightarrow S_M = L \sim d(E_1, \dots, E_n), S_D = S_Q = \text{skip} \\ \Gamma(L, E_1, \dots, E_n) = \text{GENQUANT} \rightarrow S_Q = L \sim d(E_1, \dots, E_n), S_D = S_M = \text{skip} \end{array}}{L \sim d(E_1, \dots, E_n) \Downarrow_{\Gamma} (S_D, S_M, S_Q)}$$

(SHRED IF)

$$\frac{S_1 \Downarrow_{\Gamma} (S_{D_1}, S_{M_1}, S_{Q_1}) \quad S_2 \Downarrow_{\Gamma} (S_{D_2}, S_{M_2}, S_{Q_2})}{\text{if}(g) S_1 \text{ else } S_2 \Downarrow_{\Gamma} (\text{if}(g) S_{D_1} \text{ else } S_{D_2}), (\text{if}(g) S_{M_1} \text{ else } S_{M_2}), (\text{if}(g) S_{Q_1} \text{ else } S_{Q_2})}$$

(SHRED FOR)

$$\frac{S \Downarrow_{\Gamma} (S_D, S_M, S_Q)}{\text{for}(x \text{ in } g_1 : g_2) S \Downarrow_{\Gamma} (\text{for}(x \text{ in } g_1 : g_2) S_D), (\text{for}(x \text{ in } g_1 : g_2) S_M), (\text{for}(x \text{ in } g_1 : g_2) S_Q)}$$

Here, $\Gamma(E)$ (Definition 14) gives the principal type of an expression E , while $\Gamma(E_1, \dots, E_n)$ (Definition 15) gives the least upper bound of the principal types of E_1, \dots, E_n .

The (SHRED IF) and (SHRED FOR) rules make sure to shred if and for statements so they are separated into parts that can be computed independently at each of the three levels. Note that the usage of **if** and **for** guards is simplified to avoid stating rules for when the guard(s) are of different levels. For example, if we have a statement **if**(E) S_1 **else** S_2 , where E is of level **MODEL**, then we cannot access E at level **DATA**, thus the actual shredding rule we would use is:

(SHRED IF MODEL LEVEL)

$$\frac{S_1 \Downarrow_{\Gamma} (S_{D_1}, S_{M_1}, S_{Q_1}) \quad S_2 \Downarrow_{\Gamma} (S_{D_2}, S_{M_2}, S_{Q_2})}{\text{if}(g) S_1 \text{ else } S_2 \Downarrow_{\Gamma} \text{skip}, (\text{if}(g) S_{D_1}; S_{M_1} \text{ else } S_{D_2}; S_{M_2}), (\text{if}(g) S_{Q_1} \text{ else } S_{Q_2})}$$

These shredding rules follow very closely those given by Gorinova et al. [2019]. The main difference is that sample statements ($L \sim d(E_1, \dots, E_n)$) are allowed to be of **GENQUANT** level and can be included in the last, generative slice of the program (see rule (SHRED SAMPLE)). In other words, such **GENQUANT** sample statements are those statements that can be interpreted as probabilistic assignment (using random number generator functions) to directly sample from the posterior distribution according to ancestral sampling.

We provide proofs for the following key results in Appendix A: Shredding produces *single-level statements* (Definition 5 and Lemma 4) and shredding is semantics preserving (Lemma 6).

Intuitively, a single-level statement of level ℓ is one that updates only variables of level ℓ .

Definition 5 (Single-level Statement $\Gamma \vdash \ell(S)$). We define single-level statements S of level ℓ with respect to Γ (written $\Gamma \vdash \ell(S)$), by induction:

Single Level Statements:

(ASSIGN SINGLE) $\Gamma(x) = (_, \ell)$	(SEQ SINGLE) $\Gamma \vdash \ell(S) \quad \Gamma \vdash \ell(S')$	(FOR SINGLE) $\Gamma, x : (\text{int}, \ell) \vdash \ell(S)$	(IF SINGLE) $\Gamma \vdash \ell(S_1) \quad \Gamma \vdash \ell(S_2)$
$\Gamma \vdash \ell(x[E_1] \cdots [E_n] = E)$	$\Gamma \vdash \ell(S; S')$	$\Gamma \vdash \ell(\text{for}(x \text{ in } E_1 : E_2) S)$	$\Gamma \vdash \ell(\text{if}(E) S_1 \text{ else } S_2)$
(SKIP SINGLE) $\Gamma \vdash \ell(\text{skip})$	(FACTOR SINGLE) $\Gamma \vdash E : \ell \quad \forall \ell' < \ell. \Gamma \not\vdash E : \ell'$ $\Gamma \vdash \ell(\text{factor}(E))$		

$$\frac{\text{(SAMPLE SINGLE)} \\ \Gamma \vdash L \sim d(E_1, \dots, E_n) : \ell \quad \forall \ell' < \ell. \Gamma \not\vdash L \sim d(E_1, \dots, E_n) : \ell'}{\Gamma \vdash \ell(L \sim d(E_1, \dots, E_n))}$$

LEMMA 4 (SHREDDING PRODUCES SINGLE-LEVEL STATEMENTS).

$$\Gamma \vdash S \wedge S \Downarrow_{\Gamma} (S_D, S_M, S_Q) \implies \Gamma \vdash \text{DATA}(S_D) \wedge \Gamma \vdash \text{MODEL}(S_M) \wedge \Gamma \vdash \text{GENQUANT}(S_Q).$$

We prove a result about the effect of single-level statements on the store and weight of well-typed programs (Lemma 5). Intuitively, this result shows that a single-level statement of level ℓ acts on the state and weight in a way that is independent of levels greater than ℓ .

LEMMA 5 (PROPERTY OF SINGLE-LEVEL STATEMENTS). *Let $\Gamma_{\sigma}, \Gamma_{\mathbf{x}}, S$ be a SlicStan program, such that S is a single-level statement of level ℓ , $\Gamma \vdash \ell(S)$. Then there exist unique functions f and ϕ , such that for any $\sigma, \mathbf{x} \models \Gamma_{\sigma}, \Gamma_{\mathbf{x}}$:*

$$\llbracket S \rrbracket(\sigma)(\mathbf{x}) = f(\sigma_{\leq \ell}, \mathbf{x}_{\leq \ell}) \cup \sigma_{> \ell}, \quad \phi(\sigma_{\leq \ell})(\mathbf{x}_{\leq \ell}),$$

where we write $\sigma_{\leq \ell} = \{(x \mapsto V) \in \sigma \mid \Gamma_{\sigma}(x) = (_, \ell)\}$ and $\sigma_{> \ell} = \sigma \setminus \sigma_{\leq \ell}$.

LEMMA 6 (SEMANTIC PRESERVATION OF \Downarrow_{Γ}). *If $\Gamma \vdash S : \text{DATA}$ and $S \Downarrow_{\Gamma} (S_D, S_M, S_Q)$, then $\llbracket S \rrbracket = \llbracket S_D; S_M; S_Q \rrbracket$.*

2.6 Density Factorisation

As an extension of Gorinova et al. [2019], we show that shredding induces a natural factorization of the density implemented by the program: $p(\mathcal{D}, \theta, Q) = p(\theta, \mathcal{D})p(Q \mid \theta, \mathcal{D})$.⁹ This means that we can separate the program into a deterministic preprocessing part, a part that uses a “heavy-weight” inference algorithm, such as HMC, and a part that uses simple ancestral sampling.

THEOREM 1 (SHREDDING INDUCES A FACTORISATION OF THE DENSITY). *Suppose $\Gamma \vdash S : \text{DATA}$ and $S \Downarrow_{\Gamma} S_D, S_M, S_Q$ and $\Gamma = \Gamma_{\sigma}, \Gamma_{\mathcal{D}}, \Gamma_{\theta}, \Gamma_Q$. For all $\sigma, \mathcal{D}, \theta$, and Q : if $\sigma, \mathcal{D}, \theta, Q \models \Gamma_{\sigma}, \Gamma_{\mathcal{D}}, \Gamma_{\theta}, \Gamma_Q$, and $\llbracket S \rrbracket_p(\sigma)(\mathcal{D}, \theta, Q) \propto p(\mathcal{D}, \theta, Q)$ and $\tilde{W}(S_Q) = \text{dom}(\Gamma_Q)$ then:*

- (1) $\llbracket S_M \rrbracket_p(\sigma_{\mathcal{D}})(\mathcal{D}, \theta, Q) \propto p(\theta, \mathcal{D})$,
- (2) $\llbracket S_Q \rrbracket_p(\sigma_M)(\mathcal{D}, \theta, Q) = p(Q \mid \theta, \mathcal{D})$,

where $\sigma_{\mathcal{D}} = \llbracket S_D \rrbracket_s(\sigma)(\mathcal{D}, \theta, Q)$, $\sigma_M = \llbracket S_M \rrbracket_s(\sigma_{\mathcal{D}})(\mathcal{D}, \theta, Q)$, and $p(\mathcal{D}, \theta, Q) = p(\mathcal{D}, \theta)p(Q \mid \mathcal{D}, \theta)$.

PROOF. This follows by proving a more general result using induction on the structure of S , Lemma 6, Lemma 2, and Lemma 4. See Appendix A for full proof. \square

The given SlicStan program S defines a joint density $p(\mathcal{D}, \theta, Q)$. By shredding, we obtain a **MODEL** block S_M that defines $p(\theta, \mathcal{D})$ and a **GENQUANT** block S_Q that defines $p(Q \mid \theta, \mathcal{D})$. Hence, inference in Stan using these blocks recovers the semantics $p(\mathcal{D}, \theta, Q)$ of the SlicStan program.

3 THEORY: CONDITIONAL INDEPENDENCE BY TYPING

This section presents the main theoretical contribution of the article: an information flow type system for conditional independence. We present a type system and show that a well-typed program in that system is guaranteed to have certain conditional independencies in its density semantics. As a reminder, determining the conditional independence relationships between variables is important, as such relationships capture a qualitative summary of the specified model and can facilitate more efficient inference. For example, in Section 4, we present an application that uses our type

⁹Here, $p(Q \mid \theta, \mathcal{D})$ denotes the *conditional probability density* of Q , given the values of θ and \mathcal{D} .


system: a semantic-preserving transformation that allows for discrete parameters to be introduced in SlicStan, which was previously not possible due to efficiency constraints.

Our aim is to optimise probabilistic programs by transforming abstract syntax trees or intermediate representations (as in the Stan compiler) that are close to abstract syntax. Hence, we seek a way to compute conditional dependencies by a type-based source analysis, rather than by explicitly constructing a separate graphical representation of the probabilistic model.

Given three disjoint sets of **random variables (RVs)** A , B , and C , we say that A is *conditionally independent* of B given C , written $A \perp\!\!\!\perp B \mid C$, if and only if their densities factorise as $p(A, B \mid C) = p(A \mid C)p(B \mid C)$. (An alternative formulation states that $A \perp\!\!\!\perp B \mid C$ if and only if $p(A, B, C) = \phi_1(A, C)\phi_2(B, C)$ for some functions ϕ_1 and ϕ_2 .) Deriving conditional independencies in the presence of a graphical model (such as a factor graph¹⁰) is straightforward, which is why some PPLs focus on building and performing inference on graphs (for example, Infer.NET [Minka et al. 2014]). However, building and manipulating a factor graph in generative PPLs (e.g., Gen [Cusumano-Towner et al. 2019], Pyro [Uber AI Labs 2017], Edward2 [Tran et al. 2018], PyMC3 [Salvatiere et al. 2016]), or imperative density-based PPLs (SlicStan, Stan) is not straightforward. Dependencies between modelled variables might be separated by various deterministic transformations, making it harder to track the information flow and—more importantly—more difficult to isolate parts of the model needed for transformations such as variable elimination. In the case of SlicStan, each program can still be thought of as specifying a factor graph *implicitly*. In this article, we focus on the problem of how to work with conditional independence information implicitly encoded in a probabilistic program, without having access to an explicit factor graph. For example, consider Program A:

A. Simple Hidden Markov Model (HMM)

```
int<2> z1 ~ bern( $\theta_0$ );
real  $\theta_1$  = foo( $\theta_0$ , z1);
int<2> z2 ~ bern( $\theta_1$ );
real  $\phi_1$  = foo(1, z1);
real  $\phi_2$  = foo(1, z2);
int<2> y1 ~ bern( $\phi_1$ );
int<2> y2 ~ bern( $\phi_2$ );
```



The factor graph above represents the factorisation of the joint density function over the parameters of the program: $p(z_1, z_2, y_1, y_2) = b(z_1 \mid \theta_0)b(y_1 \mid \text{foo}(1, z_1))b(z_2 \mid \text{foo}(\theta_0, z_1))b(y_2 \mid \text{foo}(1, z_2))$. Each of the four factors is represented by a square node in the graph, and it connects to the variables (circle nodes) that the factor depends on. This representation is useful for thinking about conditional independencies. For example, it is immediately evident from the graph that variables that connect to the same square node cannot be conditionally independent, as they share a factor. More generally, if there is an (uninterrupted by observed variables) direct path between two variables, then these two variables are *not* conditionally independent [Frey 2002].


When looking at the factor graph, it is straightforward to see that z_1 and z_2 are *not* conditionally independent, and neither are z_1 and y_1 nor z_2 and y_2 , as there is a direct path between each of these pairs. When looking at the program, however, we need to reason about the information flow through the deterministic variables θ_1 , ϕ_1 and ϕ_2 to reach the same conclusion.

¹⁰A factor graph is a bipartite graph that shows the factorisation of a multivariable function. Variables are circular nodes, and each factor of the function is a square node. An edge exists between a variable node x and a factor node ϕ if and only if ϕ is a function of x . See Program A and its corresponding factor graph as an example, or Koller and Friedman [2009] for details.

Moreover, manipulation of the program based on conditional dependencies can also be more difficult without a factor graph. As an example, consider the problem of variable elimination (which we discuss in more details in Section 4.3). If we are to eliminate z_1 in the factor graph, using variable elimination, we would simply merge the factors directly connected to z_1 , sum over z_1 and attach the new factors to all former neighbours of z_1 (in this case y_1 and z_2 , but not y_2). However, in the case of an imperative program, we need to isolate all the statements that depend on z_1 and group them together without changing the meaning of the program beyond the elimination:

B. HMM with z_1 marginalised out

```
factor(sum([target(
  z1 ~ bern(theta_0); real theta_1 = foo(theta_0, z1);
  z2 ~ bern(theta_1); real phi_1 = foo(1, z1);
  y1 ~ bern(phi_1); ) | z1 in 1 : 2 ]));
real phi_2 = foo(1, z2);
int<2> y2 ~ bern(phi_2);
```



We need a way to analyse the information flow to determine conditional independencies between variables. In the example above, we can leave y_2 out of the elimination of z_1 , because z_1 and y_2 are conditionally independent given z_2 , written $z_1 \perp\!\!\!\perp y_2 \mid z_2$.

To analyse the information flow, we introduce a novel type system, which we refer to via the relation \vdash_2 . It works with a lower semi-lattice $(\{\mathbf{L1}, \mathbf{L2}, \mathbf{L3}\}, \leq)$ of levels, where $\mathbf{L1} \leq \mathbf{L2}$ and $\mathbf{L1} \leq \mathbf{L3}$ and $\mathbf{L2}$ and $\mathbf{L3}$ are unrelated. (Recall that a lower semi-lattice is a partial order in which any two elements ℓ_1, ℓ_2 have a greatest lower bound $\ell_1 \sqcap \ell_2$ but do not always have an upper bound.) A well-typed program induces a conditional independence relationship for the **(random) variables (RVs)** in the program: $\mathbf{L2}$ -RVs $\perp\!\!\!\perp$ $\mathbf{L3}$ -RVs \mid $\mathbf{L1}$ -RVs.

In the example above, this result allows us to eliminate $\mathbf{L2}$ -variables (z_1), while only considering $\mathbf{L1}$ -variables (y_1 and z_2) and knowing $\mathbf{L3}$ -variables (y_2) are unaffected by the elimination. We can use a shredding relation almost identical to that of Section 2.5 to slice the program in a semantics-preserving way and isolate the sub-statements needed for elimination. Here, θ_1 and ϕ_1 must be of level $\mathbf{L2}$ for the program to be well-typed. Thus, all statements involving z_1, θ_1 , or ϕ_1 are of level $\mathbf{L2}$, and the shredding relation groups them together inside of the elimination loop for z_1 .

Figure 2 shows the relationship between the levels $\mathbf{L1}, \mathbf{L2}, \mathbf{L3}$ and the shredding relation. Information flows from $\mathbf{L1}$ to $\mathbf{L2}$ and $\mathbf{L3}$, but there is no flow of information between $\mathbf{L2}$ and $\mathbf{L3}$ (Figure 2(b)). A \vdash_2 -well-typed program S is shredded by \uparrow_{Γ} into S_1, S_2 , and S_3 , where S_1 only mentions $\mathbf{L1}$ variables, S_2 only mentions $\mathbf{L1}$ and $\mathbf{L2}$ variables, and S_3 only mentions $\mathbf{L1}$ and $\mathbf{L3}$ variables. This can be understood as a new factor graph formulation of the original program S , where each of the substatements S_1, S_2, S_3 defines a factor connected to any involved variables (Figure 2(a)).

Our approach relies on determining the $\mathbf{L1}, \mathbf{L2}, \mathbf{L3}$ level types by type inference, as they are *not* intrinsic to the variables or program in any way, but are designed solely to determine conditional independence relationships. These types are not accessible by the probabilistic programming user. Our type system makes it possible to answer various questions about conditional independence in a program. Assuming a program defining a joint density $p(\mathbf{x})$, we can use the type system to:

- (1) Check if $\mathbf{x}_2 \perp\!\!\!\perp \mathbf{x}_3 \mid \mathbf{x}_1$ for some partitioning $\mathbf{x} = \mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3$.
- (2) Find an optimal variable partitioning. Given a variable $x \in \mathbf{x}$, find a partitioning $\mathbf{x} = \mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3$, such that $x \in \mathbf{x}_2, \mathbf{x}_2 \perp\!\!\!\perp \mathbf{x}_3 \mid \mathbf{x}_1$, and \mathbf{x}_1 and \mathbf{x}_2 are as small as possible.




Fig. 2. Intuition for the semi-lattice case $\mathbf{L1} < \mathbf{L2}$ and $\mathbf{L1} < \mathbf{L3}$, where x_ℓ is of level ℓ . We get $x_{\mathbf{L2}} \perp\!\!\!\perp x_{\mathbf{L3}} \mid x_{\mathbf{L1}}$.

- (3) Ask questions about the Markov boundary of a variable. Given two variables x and x' , find the partitioning $\mathbf{x} = x, \mathbf{x}_1, \mathbf{x}_2$, such that $x \perp\!\!\!\perp \mathbf{x}_1 \mid \mathbf{x}_2$ and \mathbf{x}_2 is as small as possible. Is x' in \mathbf{x}_2 ? In other words, is x' in the Markov boundary of x ?

In the rest of Section 3, we give the \vdash_2 type system (Section 3.1), state a noninterference result (Lemma 7, Lemma 8) and show that semantics is preserved when shredding \vdash_2 -well-typed programs (Lemma 10). We present the type system and transformation rules in a declarative style. The implementation relies on type inference, which we discuss in Section 4.4. We derive a result about the way shredding factorises the density defined by the program (Theorem 2). We prove a conditional independence result (Section 3.2, Theorem 3) and discuss the scope of our approach with examples (Section 3.3).

3.1 The \vdash_2 Type System

We introduce a modified version of SlicStan's type system. Once again, types T range over pairs (τ, ℓ) of a base type τ , and a level type ℓ , but levels ℓ are one of $\mathbf{L1}$, $\mathbf{L2}$, or $\mathbf{L3}$, which form a lower semi-lattice $(\{\mathbf{L1}, \mathbf{L2}, \mathbf{L3}\}, \leq)$, where $\mathbf{L1} \leq \mathbf{L2}$ and $\mathbf{L1} \leq \mathbf{L3}$. This means, for example, that an $\mathbf{L2}$ variable can depend on an $\mathbf{L1}$ variable, but an $\mathbf{L3}$ variable cannot depend on an $\mathbf{L2}$ variable, as level types $\mathbf{L2}$ and $\mathbf{L3}$ are incomparable.

The type system is a standard information flow type system, very similar to the \vdash system introduced in Section 2.2. We mark the only non-standard rules, (SAMPLE2), (FACTOR2), and (SEQ2), which also differ from those of \vdash . (SAMPLE2) and (FACTOR2) both have the same effect as an assignment to an implicit weight variable that can be of any of the three levels. (SEQ2) is a less restrictive version of (SEQ) and exactly as in Gorinova et al. [2019], and it makes sure the program can be sliced later.

Note also the non-interference between $\mathbf{L2}$ and $\mathbf{L3}$ relies on the (PRIMCALL2) rule not being derivable when the least upper bound $\bigsqcup_{i=1}^n \ell_i$ does not exist.

Typing Rules for Expressions:

(ESUB2)	(VAR2)	(CONST2)	(ARR2)
$\frac{\Gamma \vdash_2 E : (\tau, \ell) \quad \ell \leq \ell'}{\Gamma \vdash_2 E : (\tau, \ell')}$	$\frac{}{\Gamma, x : T \vdash_2 x : T}$	$\frac{\mathbf{ty}(c) = \tau}{\Gamma \vdash_2 c : (\tau, \mathbf{L1})}$	$\frac{\Gamma \vdash_2 E_i : (\tau, \ell) \quad \forall i \in 1..n}{\Gamma \vdash_2 [E_1, \dots, E_n] : (\tau[n], \ell)}$
(ARREL2)	(PRIMCALL2)($f : \tau_1, \dots, \tau_n \rightarrow \tau$)		
$\frac{\Gamma \vdash_2 E_1 : (\tau[n], \ell) \quad \Gamma \vdash_2 E_2 : (\mathbf{int}, \ell)}{\Gamma \vdash_2 E_1[E_2] : (\tau, \ell)}$	$\frac{\Gamma \vdash_2 E_i : (\tau_i, \ell_i) \quad \forall i \in 1..n}{\Gamma \vdash_2 f(E_1, \dots, E_n) : (\tau, \bigsqcup_{i=1}^n \ell_i)}$		

$$\frac{(\text{ARRCOMP2}) \quad \forall i = 1, 2. \Gamma \vdash_2 E_i : (\mathbf{int}, \ell) \quad \Gamma, x : (\mathbf{int}, \ell) \vdash E : (\tau, \ell) \quad x \notin \text{dom}(\Gamma) \quad \Gamma \vdash_2 S : \ell'' \quad \forall \ell' > \ell. R_{\Gamma \vdash \ell'}(S) = \emptyset}{\Gamma \vdash_2 [E \mid x \mathbf{in} E_1 : E_2] : (\tau[n], \ell)} \quad \frac{(\text{TARGET2})}{\Gamma \vdash_2 \mathbf{target}(S) : (\mathbf{real}, \ell)}$$

Typing Rules for Statements:

$$\frac{(\text{SSUB2}) \quad \Gamma \vdash_2 S : \ell' \quad \ell \leq \ell'}{\Gamma \vdash_2 S : \ell} \quad \frac{(\text{ASSIGN2}) \quad \Gamma(L) = (\tau, \ell) \quad \Gamma \vdash_2 E : (\tau, \ell)}{\Gamma \vdash_2 (L = E) : \ell}$$

$\frac{(\text{SAMPLE2}) \quad \Gamma \vdash_2 \mathbf{factor}(D(L \mid E_1, \dots, E_n)) : \ell}{\Gamma \vdash_2 L \sim D_{\text{dist}}(E_1, \dots, E_n) : \ell}$	$\frac{(\text{FACTOR2}) \quad \Gamma \vdash_2 E : (\mathbf{real}, \ell)}{\Gamma \vdash_2 \mathbf{factor}(E) : \ell}$	$\frac{(\text{SEQ2}) \quad \Gamma \vdash_2 S_1 : \ell \quad \Gamma \vdash_2 S_2 : \ell \quad \mathcal{S}(S_1, S_2)}{\Gamma \vdash_2 (S_1; S_2) : \ell}$
---	--	---

$$\frac{(\text{IF2}) \quad \Gamma \vdash_2 E : (\mathbf{bool}, \ell) \quad \Gamma \vdash_2 S_1 : \ell \quad \Gamma \vdash_2 S_2 : \ell}{\Gamma \vdash_2 \mathbf{if}(E) S_1 \mathbf{else} S_2 : \ell} \quad \frac{(\text{SKIP2})}{\Gamma \vdash_2 \mathbf{skip} : \ell}$$

$$\frac{(\text{FOR2}) \quad \Gamma \vdash_2 E_1 : (\mathbf{int}, \ell) \quad \Gamma \vdash_2 E_2 : (\mathbf{int}, \ell) \quad \Gamma, x : (\mathbf{int}, \ell) \vdash_2 S : \ell \quad x \notin \text{dom}(\Gamma) \quad x \notin W(S)}{\Gamma \vdash_2 \mathbf{for}(x \mathbf{in} E_1 : E_2) S : \ell}$$

We state and prove a noninterference result for \vdash_2 , which follows similarly to the result for \vdash .

LEMMA 7 (NONINTERFERENCE OF \vdash_2). *Suppose $s_1 \models \Gamma$, $s_2 \models \Gamma$, and $s_1 \approx_\ell s_2$ for some ℓ . Then for a SlicStan statement S and expression E :*

- (1) If $\Gamma \vdash_2 E : (\tau, \ell)$ and $(s_1, E) \Downarrow V_1$ and $(s_2, E) \Downarrow V_2$, then $V_1 = V_2$.
- (2) If $\Gamma \vdash_2 S : \ell$ and $(s_1, S) \Downarrow s'_1, w_1$ and $(s_2, S) \Downarrow s'_2, w_2$, then $s'_1 \approx_\ell s'_2$.

PROOF. (1) follows by rule induction on the derivation $\Gamma \vdash_2 E : (\tau, \ell)$, and using that if $\Gamma \vdash_2 E : (\tau, \ell)$, $x \in R(E)$ and $\Gamma(x) = (\tau', \ell')$, then $\ell' \leq \ell$. (2) follows by rule induction on the derivation $\Gamma \vdash_2 S : \ell$ and using (1). \square

Once again, we derive a more convenient form of the noninterference result. Because the level types **L2** and **L3** are not comparable in the order \leq , changes in the store at **L2** do not affect the store at **L3** and vice versa.

LEMMA 8 (NONINTERFERENCE OF \vdash_2 -WELL-TYPED PROGRAMS). *Let $\Gamma_\sigma, \Gamma_x, S$ be a SlicStan program, and $\Gamma \vdash_2 S : \mathbf{L1}$. There exist unique functions f, g , and h , such that for all $\sigma \models \Gamma_\sigma, \mathbf{x} \models \Gamma_x$ and σ' such that $\llbracket S \rrbracket_s(\sigma)(\mathbf{x}) = \sigma'$:*

$$\sigma'_{\mathbf{L1}} = f(\sigma_{\mathbf{L1}}, \mathbf{x}_{\mathbf{L1}}), \quad \sigma'_{\mathbf{L2}} = g(\sigma_{\mathbf{L1}}, \sigma_{\mathbf{L2}}, \mathbf{x}_{\mathbf{L1}}, \mathbf{x}_{\mathbf{L2}}), \quad \sigma'_{\mathbf{L3}} = h(\sigma_{\mathbf{L1}}, \sigma_{\mathbf{L3}}, \mathbf{x}_{\mathbf{L1}}, \mathbf{x}_{\mathbf{L3}}).$$

PROOF. Follows from noninterference (Lemma 7). \square

Next, we extend the shredding relation from Section 2.5, and the concept of single-level statements, to SlicStan programs that are well-typed with respect to \vdash_2 . This is done by simply treating **L1** as **DATA**, **L2** as **MODEL**, and **L3** as **GENQUANT** for the purpose of shredding. We include the full

definition of shredding with respect to \vdash_2 for completeness below. We use the same notation \Downarrow_Γ , and we generally treat the *standard shredding relation* from Section 2.5 and the *conditional independence shredding relation* presented here, as the same relation, as there is no difference between the two, other than the naming of levels.

Shredding Rules for Statements:

(SHRED2 ASSIGN)	
$\Gamma(L) = \mathbf{L1} \rightarrow S_1 = L = E, S_2 = S_3 = \mathbf{skip}$	(SHRED2 SEQ)
$\Gamma(L) = \mathbf{L2} \rightarrow S_2 = L = E, S_1 = S_3 = \mathbf{skip}$	$\frac{S_1 \Downarrow_\Gamma S_1^{(1)}, S_2^{(1)}, S_3^{(1)} \quad S_2 \Downarrow_\Gamma S_1^{(2)}, S_2^{(2)}, S_3^{(2)}}{S_1; S_2 \Downarrow_\Gamma (S_1^{(1)}; S_1^{(2)}), (S_2^{(1)}; S_2^{(2)}), (S_3^{(1)}; S_3^{(2)})}$
$\Gamma(L) = \mathbf{L3} \rightarrow S_3 = L = E, S_1 = S_2 = \mathbf{skip}$	$\frac{}{L = E \Downarrow_\Gamma (S_1, S_2, S_3)}$
(SHRED2 FACTOR)	
$\Gamma(E) = \mathbf{L1} \rightarrow S_1 = \mathbf{factor}(E), S_2 = S_3 = \mathbf{skip}$	(SHRED2 SKIP)
$\Gamma(E) = \mathbf{L2} \rightarrow S_2 = \mathbf{factor}(E), S_1 = S_3 = \mathbf{skip}$	$\frac{}{\mathbf{skip} \Downarrow_\Gamma (\mathbf{skip}, \mathbf{skip}, \mathbf{skip})}$
$\Gamma(E) = \mathbf{L3} \rightarrow S_3 = \mathbf{factor}(E), S_1 = S_2 = \mathbf{skip}$	$\frac{}{\mathbf{factor}(E) \Downarrow_\Gamma (S_1, S_2, S_3)}$
(SHRED2 SAMPLE)	
$\Gamma(L, E_1, \dots, E_n) = \mathbf{L1} \rightarrow S_1 = L \sim d(E_1, \dots, E_n), S_2 = S_3 = \mathbf{skip}$	$\Gamma(L, E_1, \dots, E_n) = \mathbf{L2} \rightarrow S_2 = L \sim d(E_1, \dots, E_n), S_1 = S_3 = \mathbf{skip}$
$\Gamma(L, E_1, \dots, E_n) = \mathbf{L3} \rightarrow S_3 = L \sim d(E_1, \dots, E_n), S_1 = S_2 = \mathbf{skip}$	$\frac{}{L \sim d(E_1, \dots, E_n) \Downarrow_\Gamma (S_1, S_2, S_3)}$
(SHRED2 IF)	
$\mathbf{if}(g) S_1 \mathbf{else} S_2 \Downarrow_\Gamma (\mathbf{if}(g) S_1^{(1)} \mathbf{else} S_1^{(2)}), (\mathbf{if}(g) S_2^{(1)} \mathbf{else} S_2^{(2)}), (\mathbf{if}(g) S_3^{(1)} \mathbf{else} S_3^{(2)})$	$\frac{S_1 \Downarrow_\Gamma S_1^{(1)}, S_2^{(1)}, S_3^{(1)} \quad S_2 \Downarrow_\Gamma S_1^{(2)}, S_2^{(2)}, S_3^{(2)}}{\mathbf{if}(g) S_1 \mathbf{else} S_2 \Downarrow_\Gamma (\mathbf{if}(g) S_1^{(1)} \mathbf{else} S_1^{(2)}), (\mathbf{if}(g) S_2^{(1)} \mathbf{else} S_2^{(2)}), (\mathbf{if}(g) S_3^{(1)} \mathbf{else} S_3^{(2)})}$
(SHRED2 FOR)	
$\mathbf{for}(x \mathbf{in} g_1 : g_2) S \Downarrow_\Gamma (\mathbf{for}(x \mathbf{in} g_1 : g_2) S_1), (\mathbf{for}(x \mathbf{in} g_1 : g_2) S_2), (\mathbf{for}(x \mathbf{in} g_1 : g_2) S_3)$	$\frac{S \Downarrow_\Gamma S_1, S_2, S_3}{\mathbf{for}(x \mathbf{in} g_1 : g_2) S \Downarrow_\Gamma (\mathbf{for}(x \mathbf{in} g_1 : g_2) S_1), (\mathbf{for}(x \mathbf{in} g_1 : g_2) S_2), (\mathbf{for}(x \mathbf{in} g_1 : g_2) S_3)}$

As before, shredding produces single-level statements, and shredding preserves semantics with respect to \vdash_2 -well-typed programs.

LEMMA 9 (SHREDDING PRODUCES SINGLE-LEVEL STATEMENTS, \vdash_2). *b If $S \Downarrow_\Gamma S_1, S_2, S_3$, then $\Gamma \vdash \mathbf{L1}(S_1)$, $\Gamma \vdash \mathbf{L2}(S_2)$, and $\Gamma \vdash \mathbf{L3}(S_3)$.*

LEMMA 10 (SEMANTIC PRESERVATION OF $\Downarrow_\Gamma, \vdash_2$). *If $\Gamma \vdash_2 S : \mathbf{L1}$ and $S \Downarrow_\Gamma S_1, S_2, S_3$, then $\llbracket S \rrbracket = \llbracket S_1; S_2; S_3 \rrbracket$.*

In addition, we derive a result about the effect of single-level statements on the store and weight of \vdash_2 -well-typed programs.

LEMMA 11 (PROPERTY OF \vdash_2 SINGLE-LEVEL STATEMENTS). *Let $\Gamma_\sigma, \Gamma_x, S$ be a SlicStan program, and $\Gamma \vdash_2 S : \mathbf{L1}$, and S be single-level statement of level ℓ , $\Gamma \vdash_2 \ell(S)$. Then there exist unique functions f and ϕ , such that for any $\sigma, x \models \Gamma_\sigma, \Gamma_x$:*

- (1) *If $\ell = \mathbf{L1}$, then $\llbracket S \rrbracket(\sigma)(x) = (f(\sigma_{\mathbf{L1}}, \mathbf{x}_{\mathbf{L1}}), \sigma_{\mathbf{L2}}, \sigma_{\mathbf{L3}}), \phi(\sigma_{\mathbf{L1}})(\mathbf{x}_{\mathbf{L1}})$,*
- (2) *If $\ell = \mathbf{L2}$, then $\llbracket S \rrbracket(\sigma)(x) = (\sigma_{\mathbf{L1}}, f(\sigma_{\mathbf{L1}}, \sigma_{\mathbf{L2}}, \mathbf{x}_{\mathbf{L1}}, \mathbf{x}_{\mathbf{L2}}), \sigma_{\mathbf{L3}}), \phi(\sigma_{\mathbf{L1}}, \sigma_{\mathbf{L2}})(\mathbf{x}_{\mathbf{L1}}, \mathbf{x}_{\mathbf{L2}})$,*
- (3) *If $\ell = \mathbf{L3}$, then $\llbracket S \rrbracket(\sigma)(x) = (\sigma_{\mathbf{L1}}, \sigma_{\mathbf{L2}}, f(\sigma_{\mathbf{L1}}, \sigma_{\mathbf{L3}}, \mathbf{x}_{\mathbf{L1}}, \mathbf{x}_{\mathbf{L3}})), \phi(\sigma_{\mathbf{L1}}, \sigma_{\mathbf{L3}})(\mathbf{x}_{\mathbf{L1}}, \mathbf{x}_{\mathbf{L3}})$.*

We give proofs for Lemma 9, 10, and 11 in Appendix A. These results allow us to derive the second key theorem of this article, Theorem 2, which, similarly to Theorem 1, gives us a result on the way shredding factorises the density defined by the program.

Here, and throughout the article, we use subscripts to refer to specific subsets of Γ . For example, Γ_{L1} stands for the subset of the parameters Γ_x , such that $x : (\tau, \ell) \in \Gamma_{L1}$ if and only if $x : (\tau, \ell) \in \Gamma_x$ and $\ell = L1$.

THEOREM 2 (SHREDDING INDUCES A FACTORISATION OF THE DENSITY (2)). *Suppose $\Gamma \vdash_2 S : L1$ with $\Gamma = \Gamma_\sigma, \Gamma_{L1}, \Gamma_{L2}, \Gamma_{L3}$, $S \Downarrow_\Gamma S_1, S_2, S_3$. Then for $\sigma, \theta_1, \theta_2, \theta_3 \models \Gamma_\sigma, \Gamma_1, \Gamma_2, \Gamma_3$, and σ', σ'' such that $\llbracket S_1 \rrbracket(\sigma)(\theta_1, \theta_2, \theta_3) = \sigma'$, and $\llbracket S_2 \rrbracket(\sigma')(\theta_1, \theta_2, \theta_3) = \sigma''$, we have:*

- (1) $\llbracket S_1 \rrbracket_p(\sigma)(\theta_1, \theta_2, \theta_3) = \phi_1(\theta_1)$,
- (2) $\llbracket S_2 \rrbracket_p(\sigma')(\theta_1, \theta_2, \theta_3) = \phi_2(\theta_1, \theta_2)$,
- (3) $\llbracket S_3 \rrbracket_p(\sigma'')(\theta_1, \theta_2, \theta_3) = \phi_3(\theta_1, \theta_3)$.

PROOF. By applying Lemma 11 to each of S_1, S_2, S_3 , which are single-level statements (Lemma 9). \square

3.2 Conditional Independence Result for \vdash_2 -Well-typed Programs

Theorem 3 states the key theoretical result of this article: the typing in programs well-typed with respect to \vdash_2 corresponds to a conditional independence relationship. In our proofs, we use the factorisation characterisation of conditional independence stated by Definition 6. This is a well-known result in the literature (e.g., Murphy [2012], Theorem 2.2.1).

Definition 6 (Characterisation of Conditional Independence as Factorisation). For variables x, y, z and a density $p(x, y, z)$, x is conditionally independent of y given z with respect to p , written $x \perp\!\!\!\perp_p y \mid z$, if and only if $\exists \phi_1, \phi_2$ such that $p(x, y, z) = \phi_1(x, z)\phi_2(y, z)$.

An equivalent formulation is $p(x, y \mid z) = p(x \mid z)p(y \mid z)$.

We extend the notion of conditional independence to apply to a general function $\phi(x, y, z)$, using the notation $x \perp\!\!\!\perp_\phi y \mid z$ to mean $\exists \phi_1, \phi_2$ such that $\phi(x, y, z) = \phi_1(x, z)\phi_2(y, z)$.

THEOREM 3 (\vdash_2 -WELL-TYPED PROGRAMS INDUCE A CONDITIONAL INDEPENDENCE RELATIONSHIP). *For a SlicStan program Γ, S such that $\Gamma \vdash_2 S : L1$, $\Gamma = \Gamma_\sigma, \Gamma_{L1}, \Gamma_{L2}, \Gamma_{L3}$, and for $\sigma, \theta_1, \theta_2, \theta_3 \models \Gamma_\sigma, \Gamma_{L1}, \Gamma_{L2}, \Gamma_{L3}$, we have $\theta_2 \perp\!\!\!\perp_\phi \theta_3 \mid \theta_1$.*

When $\llbracket S \rrbracket_p(\sigma)(\theta_1, \theta_2, \theta_3) \propto p(\theta_1, \theta_2, \theta_3)$, we have $\theta_2 \perp\!\!\!\perp_p \theta_3 \mid \theta_1$.

PROOF. Let $\theta = \theta_1, \theta_2, \theta_3$, $S \Downarrow_\Gamma S_1, S_2, S_3$, and let σ' and σ'' be such that $\sigma' = \llbracket S_1 \rrbracket_s(\sigma)(\theta)$, and $\sigma'' = \llbracket S_2 \rrbracket_s(\sigma')(\theta)$. Then, by semantic preservation of shredding (Lemma 10), we have


$$\begin{aligned}
 \llbracket S \rrbracket_p(\sigma)(\theta) &= \llbracket S_1; S_2; S_3 \rrbracket_p(\sigma)(\theta) && \text{by Lemma 10,} \\
 &= \llbracket S_1 \rrbracket_p(\sigma)(\theta) \times \llbracket S_2 \rrbracket_p(\sigma')(\theta) \times \llbracket S_3 \rrbracket_p(\sigma'')(\theta) && \text{by Lemma 2,} \\
 &= \phi_1(\theta_1) \times \phi_2(\theta_1, \theta_2) \times \phi_3(\theta_1, \theta_3) && \text{by Theorem 2,} \\
 &= \phi'(\theta_1, \theta_2) \times \phi_3(\theta_1, \theta_3),
 \end{aligned}$$

for some ϕ_1, ϕ_2 , and ϕ_3 , $\phi'(\theta_1, \theta_2) = \phi_1(\theta_1) \times \phi_2(\theta_1, \theta_2)$. Thus, $\theta_2 \perp\!\!\!\perp_\phi \theta_3 \mid \theta_1$ by definition of $\perp\!\!\!\perp_\phi$.

Suppose $\phi(\theta_1, \theta_2, \theta_3) \propto p(\theta_1, \theta_2, \theta_3)$. Then $p(\theta_1, \theta_2, \theta_3) = \phi(\theta_1, \theta_2, \theta_3) \times Z = \phi'(\theta_1, \theta_2) \times \phi_3(\theta_1, \theta_3) \times Z = \phi'(\theta_1, \theta_2) \times \phi''(\theta_1, \theta_3)$, where Z is a constant and $\phi''(\theta_1, \theta_3) = \phi_3(\theta_1, \theta_3) \times Z$. Therefore, $\theta_2 \perp\!\!\!\perp_p \theta_3 \mid \theta_1$. \square

C. Cross Model

real $x_1 \sim \text{normal}(\theta, 1)$
real $x_2 \sim \text{normal}(\theta, 1)$
real $x_3 \sim \text{normal}(x_1+x_2, 1)$
real $x_4 \sim \text{normal}(x_3, 1)$
real $x_5 \sim \text{normal}(x_3, 1)$



(a) A simple ‘cross’ model.

(b) Graphical model.

$x_1 \perp\!\!\!\perp x_2$	$x_1 \perp\!\!\!\perp x_4 \mid \{x_3\} \cup A, \forall A \subseteq \{x_2, x_5\}$
$x_1 \perp\!\!\!\perp x_5 \mid \{x_3\} \cup A, \forall A \subseteq \{x_2, x_4\}$	$x_2 \perp\!\!\!\perp x_4 \mid \{x_3\} \cup A, \forall A \subseteq \{x_1, x_5\}$
$x_2 \perp\!\!\!\perp x_5 \mid \{x_3\} \cup A, \forall A \subseteq \{x_1, x_4\}$	$x_4 \perp\!\!\!\perp x_5 \mid \{x_3\} \cup A, \forall A \subseteq \{x_1, x_2\}$

(c) CI relationships.

Fig. 3. The cross model, as written in SlicStan (a) with its DAG (b) and CI relationships (c).

3.3 Scope of the Conditional Independence Result

We have shown that \vdash_2 -well-typed programs exhibit a conditional independence relationship in their density semantics. However, it is *not* the case that every conditional independence relationship can be derived from the type system. In particular, we can only derive results of the form $\theta_2 \perp\!\!\!\perp \theta_3 \mid \theta_1$, where $\theta_1, \theta_2, \theta_3$ is a partitioning of $\theta \models \Gamma_x$ for a SlicStan program $\Gamma_\sigma, \Gamma_x, S$. That is, the relationship includes *all* parameters in the program.

We discuss the scope of our approach using an example and show a situation where trying to derive a conditional independence result that *does not* hold results in a failure to type check.

3.3.1 Example of \vdash_2 -well-typed Program \rightarrow Conditional Independence. Consider the Cross Model in Figure 3, its SlicStan program (a), its directed graphical model, (b) and the **conditional independence (CI)** relationships that hold for that model (c).

Out of the many relationships above, we can derive *all* relationships that involve *all* the variables. That is, we can use our type system to derive all conditional independence relationships that hold and are of the form $A \perp\!\!\!\perp B \mid C$, where A, B, C is some partitioning of $\{x_1, \dots, x_5\}$. However, note the following properties of conditional independence:

$$A \perp\!\!\!\perp B \mid C \iff B \perp\!\!\!\perp A \mid C \quad \text{and} \quad A \perp\!\!\!\perp B_1, B_2 \mid C \iff A \perp\!\!\!\perp B_1 \mid C \text{ and } A \perp\!\!\!\perp B_2 \mid C.$$

Some of the relationships above can be combined and written in other ways, e.g., $x_1 \perp\!\!\!\perp x_4 \mid x_2, x_3$ and $x_1 \perp\!\!\!\perp x_5 \mid x_2, x_3$ can be written as a single relationship $x_1 \perp\!\!\!\perp x_4, x_5 \mid x_2, x_3$, thus expressing them as a single relationship that includes all variables in the program.

Exploring different mappings between the parameters x_1, \dots, x_5 and the type levels **L1, L2, L3**, for which the above program typechecks, we can derive *all* CI relationships that hold for this model, except for one: $x_1 \perp\!\!\!\perp x_2$, which we cannot derive with our approach.

3.3.2 Conditional Independence Relationship does not hold \rightarrow Type Error. Suppose that we try to derive the result $x_1 \perp\!\!\!\perp x_2 \mid x_3, x_4, x_5$. This *does not hold* for Program C. By Theorem 3, we have that a program being \vdash_2 -well-typed implies that **L2 $\perp\!\!\!\perp$ L3 \mid L1**. So, we can derive $x_1 \perp\!\!\!\perp x_2 \mid x_3, x_4, x_5$ using Theorem 3 if we show that $\Gamma \vdash_2 S : \mathbf{L1}$, for $\Gamma = \{x_1 : \mathbf{L2}, x_2 : \mathbf{L3}, x_3 : \mathbf{L1}, x_4 : \mathbf{L1}, x_5 : \mathbf{L1}\}$ and S being Program C.

To typecheck $\Gamma \vdash_2 S : \mathbf{L1}$, we need to typecheck $x_3 \sim \text{normal}(x_1 + x_2, 1)$ at some level ℓ . Thus, by (SAMPLE2) and (PRIMCALL2), x_1 , x_2 , and x_3 need to typecheck at ℓ . The types of x_1 , x_2 , and x_3 are $\mathbf{L2}$, $\mathbf{L3}$, and $\mathbf{L1}$, respectively. So, using (ESUB2), it must be the case that $\mathbf{L2} \leq \ell$, and $\mathbf{L3} \leq \ell$, and $\mathbf{L1} \leq \ell$. However, no such level exists in our lower semi-lattice, as $\mathbf{L2}$ and $\mathbf{L3}$ have no upper bound. Therefore, typechecking fails and we cannot derive $x_1 \perp\!\!\!\perp x_2 \mid x_3, x_4, x_5$.

4 APPLICATION: DISCRETE PARAMETERS SUPPORT THROUGH A SEMANTICS-PRESERVING TRANSFORMATION

This section presents the main practical contribution of our work: a semantics-preserving procedure for transforming a probabilistic program to enable combined inference of discrete and continuous model parameters, which we have implemented for SlicStan. The procedure corresponds to **variable elimination (VE)** for discrete parameters implemented in the probabilistic program itself, which can be combined with gradient-based methods, such as HMC, to perform inference on all parameters.

PPLs that have gradient-based methods in the core of their inference strategy do not, in general, support directly working with discrete parameters. Stan disallows discrete model parameters altogether, while Pyro [Uber AI Labs 2017] and Edward2 [Tran et al. 2018] throw a runtime error whenever discrete parameters are used within a gradient-based method. However, working with discrete parameters in these languages is still possible, albeit in an implicit way. In many cases, discrete parameters can be marginalised out manually, and then drawn conditionally on the continuous parameters. Stan’s user guide shows many examples of this approach [Stan Development Team 2019a, Chapter 7]. Pyro provides an on-request marginalisation functionality, which automates this implicit treatment for plated factor graphs [Obermeyer et al. 2019].

The key idea of the workaround is to marginalise out the discrete parameters by hand, so the resulting program corresponds to a density function that does not depend on any discrete parameters. That is, the user writes a program that computes $\sum_{\theta_d} p(\theta_d, \theta_c) = p(\theta_c)$, where the density semantics of the original program was $p(\theta_d, \theta_c)$ for discrete parameters θ_d and continuous parameters θ_c . This allows for continuous parameters of the program to be sampled with HMC, or other gradient-based inference algorithms, whereas that would have not been possible for the program with both discrete and continuous latent variables.

Because a SlicStan program computes a density directly, it is easy to modify it to marginalise a variable. For a SlicStan program Γ, S , with parameters $\mathbf{x} \models \Gamma_{\mathbf{x}}$, and a discrete parameter z of type $\mathbf{int}\langle K \rangle$, the program $\mathbf{elim}(\mathbf{int}\langle K \rangle z) S \triangleq \mathbf{factor}(\text{sum}([\mathbf{target}(S) \mid z \mathbf{in} 1 : K])^{11})$ marginalises z :

$$\llbracket \mathbf{factor}(\text{sum}([\mathbf{target}(S) \mid z \mathbf{in} 1 : K]) \rrbracket_p(\sigma)(\mathbf{x}) = \sum_{z=1}^K \llbracket S \rrbracket_p(\sigma)(\mathbf{x}) \propto \sum_{z=1}^K p(\mathbf{x}) = p(\mathbf{x} \setminus \{z\}).$$

In other words, we can easily marginalise out all discrete variables in a probabilistic program by encapsulating the entire program in nested loops (nested array comprehension expressions in our examples). However, this approach becomes infeasible for more than a few variables. Variable elimination [Koller and Friedman 2009; Zhang and Poole 1994] exploits the structure of a model to do as little work as possible. Consider the HMM snippet (Program D) with three discrete (binary) hidden variables z_1 , z_2 , and z_3 , and observed outcomes y_1 , y_2 , and y_3 . Naively marginalising out the hidden variables results in nested loops around the original program (Program E). In the general case of N hidden variables, the resulting program is of complexity $O(2^N)$.

¹¹Here, we assume the function `sum` is available in the language.

D. A Hidden Markov Model (HMM)

```

...
int<2> z1 ~ bernoulli(theta[1]);
int<2> z2 ~ bernoulli(theta[z1]);
int<2> z3 ~ bernoulli(theta[z2]);
data real y1 ~ normal(phi[z1], 1);
data real y2 ~ normal(phi[z2], 1);
data real y3 ~ normal(phi[z3], 1);

```

E. Inefficient marginalisation

```

...
factor(sum [target(
  factor(sum [target(
    factor(sum [target(
      z1 ~ bernoulli(theta[1]);
      z2 ~ bernoulli(theta[z1]);
      z3 ~ bernoulli(theta[z2]);
      y1 ~ normal(phi[z1], 1);
      y2 ~ normal(phi[z2], 1);
      y3 ~ normal(phi[z3], 1);
    | z1 in 1:2]);
  | z2 in 1:2]);
| z3 in 1:2]);

```

F. Efficient marginalisation

```

...
real[2] f1 = // new factor on z2
[sum([target(
  z1 ~ bernoulli(theta[1]);
  z2 ~ bernoulli(theta[z1]);
  y1 ~ normal(phi[z1], 1);
  | z1 in 1:2])
| z2 in 1:2])

real[2] f2 = // new factor on z3
[sum([target(
  factor(f1[z2]);
  y2 ~ normal(phi[z2], 1);
  z3 ~ bernoulli(theta[z2]);
  | z2 in 1:2])
| z3 in 1:2])

factor(sum [target(
  factor(f2[z3]);
  y3 ~ normal(phi[z3], 1);
  | z3 in 1:2]);

```

However, this is wasteful: Expressions like $z_3 \sim \text{bernoulli}(\theta[z_2])$ do not depend on z_1 , and so do not need to be inside of the z_1 -elimination loop. VE avoids this problem by pre-computing some of the work. Program F implements VE for this model: When eliminating a variable, say, z_1 , we pre-compute statements that involve z_1 for each possible value of z_1 and store the resulting density contributions in a new factor, f_1 . This new factor depends on the variables involved in those statements—the neighbours of z_1 —in this case that is solely z_2 . We then repeat the procedure for the other variables, re-using the already computed factors where possible.

In the special case of an HMM, and given a suitable elimination order, variable elimination recovers the celebrated forward algorithm [Rabiner 1989], which has time complexity $O(N)$. Our goal is to automatically translate the source code of Program D to Program F, exploiting statically detectable independence properties in the model.

4.1 Goal

Our ultimate goal is to transform a program S with continuous parameters θ_c , discrete parameters θ_d , data \mathcal{D} , and density semantics $\llbracket S \rrbracket_p(\sigma)(\theta_d, \theta_c, \mathcal{D}) \propto p(\theta_d, \theta_c \mid \mathcal{D})$ into two subprograms: S_{HMC} and S_{GEN} , such that:

- The density defined by S_{HMC} is the marginal $p(\theta_c \mid \mathcal{D})$, with the discrete parameters θ_d marginalised out. This first statement, S_{HMC} , represents the marginalisation part of the program (see Section 4.3) and allows for **Hamiltonian Monte Carlo (HMC)** sampling of θ_c , as it does not involve any discrete parameters.
- The density defined by S_{GEN} is the conditional $p(\theta_d \mid \theta_c, \mathcal{D})$. This second statement, S_{GEN} , represents the generative part of the program (Section 4.5) and it encodes a way to draw θ_d generatively, without using HMC or another heavy-weight inference algorithm.

Similarly to the extended SlicStan slicing based on information-flow type inference, here we also want to transform and slice into sub-programs, each focusing on a subset of the parameters and preserving the overall meaning:

$$\llbracket S \rrbracket_p \propto p(\theta_d, \theta_c \mid \mathcal{D}) = p(\theta_c \mid \mathcal{D}) \times p(\theta_d \mid \theta_c, \mathcal{D}) \propto \llbracket S_{\text{HMC}} \rrbracket_p \times \llbracket S_{\text{GEN}} \rrbracket_p = \llbracket S_{\text{HMC}}; S_{\text{GEN}} \rrbracket_p.^{12}$$

Our approach performs a semantics-preserving transformation, guided by information-flow and type inference, which creates an efficient program-specific inference algorithm automatically, combining HMC with variable elimination.

4.2 Key Insight

The key practical insight of this work is to use the adaptation of SlicStan’s level types of Section 3 and its information flow type system to rearrange the program in a semantics-preserving way, so discrete parameters can be forward-sampled, instead of sampled using a heavy-weight inference algorithm. We achieve this by a program transformation for each of the discrete variables. Assuming that we are applying the transformation with respect to a variable z , we use:

- The *top-level information flow type system* $\Gamma \vdash S : \text{DATA}$ from Section 2.2, which involves the level types $\text{DATA} \leq \text{MODEL} \leq \text{GENQUANT}$. This partitions the modelled variables x into data \mathcal{D} , model parameters θ , and generated quantities Q . When we use type inference for \vdash in conjunction with shredding $S \Downarrow_{\Gamma} S_D, S_M, S_Q$ (Section 2.5), we slice the statement S into a data part S_D (involving only variables in \mathcal{D}), a non-generative part S_M (involving \mathcal{D} and θ), and a generative part S_Q (involving \mathcal{D} , θ , and Q).
- The *conditional independence information flow type system*, $\Gamma \vdash_2 S : \text{L1}$ from Section 3, which uses a lower semi-lattice of level types $\text{L1} \leq \text{L2}, \text{L1} \leq \text{L3}$. A \vdash_2 -well-typed program induces a conditional independence relationship: L2 -variables are conditionally independent of L3 -variables given L1 -variables: $x_{\text{L2}} \perp\!\!\!\perp x_{\text{L3}} \mid x_{\text{L1}}$, where $x = x_{\text{L1}}, x_{\text{L2}}, x_{\text{L3}} = \theta, \mathcal{D}$. When we use type inference for \vdash_2 in conjunction with shredding $S \Downarrow_{\Gamma} S_1, S_2, S_3$ (Section 2.5), we isolate S_2 : a part of the program that does not interfere with S_3 . We can marginalise out L2 -variables in that sub-statement only, keeping the rest of the program unchanged.
- The *discrete variable transformation relation* $\Gamma, S \xrightarrow{z} \Gamma', S'$ (defined in Section 4.6.2), which takes a SlicStan program Γ, S that has discrete model parameter z , and transforms it to a SlicStan program Γ', S' , where z is no longer a MODEL -level parameter but instead one of level GENQUANT . We define the relation in terms of \vdash and \vdash_2 as per the (ELIM GEN) rule.


4.3 Variable Elimination

VE [Koller and Friedman 2009; Zhang and Poole 1994] is an exact inference algorithm often phrased in terms of factor graphs. It can be used to compute prior or posterior marginal distributions by eliminating, one-by-one, variables that are irrelevant to the distribution of interest. VE uses dynamic programming combined with a clever use of the distributive law of multiplication over addition to efficiently compute a nested sum of a product of expressions.


We already saw an example of variable elimination in Section 3 (Programs A and B). The idea is to eliminate (marginalise out) variables one-by-one. To eliminate a variable z , we multiply all of the factors connected to z to form a single expression, then sum over all possible values for z to create a new factor, remove z from the graph, and finally connect the new factor to all former neighbours¹³ of z . Recall Program D, with latent variables z_1, z_2, z_3 and observed data $y = y_1, y_2, y_3$.

¹²This expression is simplified for readability.


¹³“Neighbours” refers to the variables that are connected to a factor that connects to z .




(a) To eliminate z_1 , we remove z_1 and all its neighbouring factors (in red). Create a new factor f_1 , by summing out z_1 from the product of these factors.



(c) To eliminate z_2 , we remove z_2 and all its neighbouring factors (in red). Create a new factor f_2 , by summing out z_2 from the product of these factors.



(b) Connect f_1 (in green) to the former neighbours of z_1 . The remaining factor graph defines the marginal $p(z_2, z_3 | y)$.



(d) Connect f_2 (in green) to the former neighbours of z_2 . The remaining factor graph defines the marginal $p(z_3 | y)$.

Fig. 4. Step-by-step example of variable elimination.

Figure 4 shows the VE algorithm step-by-step applied to this program. We eliminate z_1 to get the marginal on z_2 and z_3 (Figure 4(a) and (b)), then eliminate z_2 to get the marginal on z_3 (Figure 4(c) and (d)).

4.4 Conditional Independence Relationships and Inferring the Markov Blanket

The key property we are looking for to be able to marginalise out a variable independently of another is conditional independence given neighbouring variables. If we shred a \vdash_2 -well-typed program into S_1, S_2 , and S_3 , and think of $\llbracket S_1 \rrbracket_p$, $\llbracket S_2 \rrbracket_p$, and $\llbracket S_3 \rrbracket_p$ as factors, it is easy to visualise the factor graph corresponding to the program: it is as in Figure 5(a). Eliminating all x_{L2} variables, ends up only modifying the $\llbracket S_2 \rrbracket_p$ factor (Figure 5(b)).

When using VE to marginalise out a parameter z , we want to find the smallest set of other parameters A , such that $z \perp\!\!\!\perp B | A$, where B is the rest of the parameters. The set A is also called z 's minimal *Markov blanket* or *Markov boundary*. Once we know this set, we can ensure that we involve the smallest possible number of variables in z 's elimination, which is important to achieve a performant algorithm.

For example, when we eliminate z_1 in Program D, both z_2 and y_1 need to be involved, as z_1 shares a factor with them. By contrast, there is no need to include y_2, z_3, y_3 and the statements associated with them, as they are unaffected by z_1 , given z_2 . The variables y_1 and z_2 form z_1 's Markov blanket: Given these variables, z_1 is conditionally independent of all other variables. That is, $z_1 \perp\!\!\!\perp z_3, y_2, y_3 | z_2, y_1$.

The type system we present in Section 3 can tell us if the conditional independence relationship $x_{L2} \perp\!\!\!\perp x_{L3} | x_{L1}$ holds for a concrete partitioning of the modelled variables $x = x_{L1}, x_{L2}, x_{L3}$. But to find the Markov blanket of a variable z we want to eliminate, we rely on type inference. We define a performance ordering between the level types $L3 < L1 < L2$, where our first preference is for variables to be of level $L3$, level $L1$ is our second preference, and $L2$ is our last resort. In our implementation, we use bidirectional type-checking [Pierce and Turner 2000] to synthesise hard constraints imposed by the type system, and resolve them, while optimising for the soft



(a) A \vdash_2 -well-typed program with parameters \mathbf{x} . (b) Eliminating \mathbf{x}_{L2} consists of modifying only $\llbracket S_2 \rrbracket_p$.

Fig. 5. The factor graph and VE induced by the shedding $S \uparrow_{\Gamma} S_1, S_2, S_3$ according to the semi-lattice $\mathbf{L1} \leq \mathbf{L2}, \mathbf{L3}$.

constraints given by the $<$ ordering. This maximises the number of variables that are conditionally independent of z given its blanket ($\mathbf{L3}$) and minimises the number of variables forming the blanket ($\mathbf{L1}$). Fixing z to be of $\mathbf{L2}$ level, and $\mathbf{L2}$ being the least preferred option, ensures that only z and variables dependent on z through deterministic assignment are of that level.

4.5 Sampling the Discrete Parameters

Variable elimination gives a way to efficiently marginalise out a variable z from a model defining density $p(\mathbf{x})$, to obtain a new density $p(\mathbf{x} \setminus \{z\})$. In the context of SlicStan, this means we have the tools to eliminate all discrete parameters θ_d , from a density $p(\mathcal{D}, \theta_c, \theta_d)$ on data \mathcal{D} , continuous parameters θ_c and discrete parameters θ_d . The resulting marginal $\sum_{\theta_d} p(\mathcal{D}, \theta_c, \theta_d) = p(\mathcal{D}, \theta_c)$ does not involve discrete parameters and, therefore, we can use gradient-based methods to infer θ_c . However, the method so far does not give us a way to infer the discrete parameters θ_d .

To infer these, we observe that $p(\mathbf{x}) = p(\mathbf{x} \setminus \{z\})p(z \mid \mathbf{x} \setminus \{z\})$, which means that we can preserve the semantics of the original model (which defines $p(\mathbf{x})$), by finding an expression for the conditional $p(z \mid \mathbf{x} \setminus \{z\})$. If $\mathbf{x}_1, \mathbf{x}_2$ is a partitioning of $\mathbf{x} \setminus \{z\}$ such that $z \perp\!\!\!\perp \mathbf{x}_2 \mid \mathbf{x}_1$, then (from Definition 6) $p(\mathbf{x}) = \phi_1(z, \mathbf{x}_1)\phi_2(\mathbf{x}_1, \mathbf{x}_2)$ for some functions ϕ_1 and ϕ_2 . Thus, $p(z \mid \mathbf{x} \setminus \{z\}) = \phi_1(z, \mathbf{x}_1) \cdot (\phi_2(\mathbf{x}_1, \mathbf{x}_2)/p(\mathbf{x} \setminus \{z\})) \propto \phi_1(z, \mathbf{x}_1)$.

In the case when z is a discrete variable of finite support, we can calculate the conditional probability exactly: $p(z \mid \mathbf{x} \setminus \{z\}) = \frac{\phi_1(z, \mathbf{x}_1)}{\sum_z \phi_1(z, \mathbf{x}_1)}$. We can apply this calculation to the factorisation of a program $\Gamma \vdash_2 S$ that is induced by shredding (Theorem 2). In that case, $\mathbf{x}_{L2}, \mathbf{x}_{L1}, \llbracket S_2 \rrbracket_p$ play the roles of z, \mathbf{x}_1 , and ϕ_1 , respectively. Consequently, we obtain a formula for drawing \mathbf{x}_{L2} conditional on the other parameters: $\mathbf{x}_{L2} \sim \text{categorical}(\llbracket S_2 \rrbracket_p(\mathbf{x}_{L2}, \mathbf{x}_{L1}) \mid \mathbf{x}_{L2} \in \text{supp}(\mathbf{x}_{L2}))$.

4.6 A Semantics-preserving Transformation Rule

In this section, we define a source-to-source transformation that implements a single step of variable elimination. The transformation re-writes a SlicStan program Γ, S with a discrete **MODEL**-level parameter z , to a SlicStan program, where z is a **GENQUANT**-level parameter. Combining the rule with the shredding presented in Section 2 results in support for efficient inference (see Section 4.8 for discussion of limitations) of both discrete and continuous random variables, where continuous variables can be inferred using gradient-based methods, such as HMC or variational inference, while discrete variables are generated using ancestral sampling. The transformation allows for SlicStan programs with explicit use of discrete parameters to be translated to Stan. We show a step-by-step example of our discrete parameter transformation in Section 4.7.

4.6.1 The ϕ , *elim*, and *gen* derived forms. We introduce three derived forms that allow us to state the rule concisely.

Variable Elimination Derived Forms

$$\begin{aligned}
\mathbf{elim}(\mathbf{int}\langle K \rangle z) S &\triangleq \mathbf{factor}(\text{sum}([\mathbf{target}(S) \mid z \mathbf{in} 1 : K])) \\
\phi(\mathbf{int}\langle K_1 \rangle z_1, \dots, \mathbf{int}\langle K_N \rangle z_N) S & \\
&\triangleq [\dots [\mathbf{target}(S) \mid z_1 \mathbf{in} 1 : K_1] \mid \dots \mid z_N \mathbf{in} 1 : K_N] \\
\mathbf{gen}(\mathbf{int}\langle K \rangle z) S &\triangleq z \sim \text{categorical}([\mathbf{target}(S) \mid z \mathbf{in} 1 : K])
\end{aligned}$$

The elimination expression $\mathbf{elim}(\mathbf{int}\langle K \rangle z) S$ adds a new factor that is equivalent to marginalising z in S . In other words, $\llbracket \mathbf{elim}(\mathbf{int}\langle K \rangle z) S \rrbracket_p(\sigma)(\mathbf{x}) = \sum_{z=1}^K \llbracket S \rrbracket_p(\sigma)(\mathbf{x})$ (see Lemma 14). A ϕ -expression $\phi(\mathbf{int}\langle K_1 \rangle z_1, \dots, \mathbf{int}\langle K_N \rangle z_N) S$ simply computes the density of the statement S in a multidimensional array for all possible values of the variables z_1, \dots, z_N . In other words, $\llbracket (f = \phi(\mathbf{int}\langle K_1 \rangle z_1, \dots, \mathbf{int}\langle K_N \rangle z_N) S); \mathbf{factor}(f[z_1] \dots [z_N]) \rrbracket_p(\sigma)(\mathbf{x}) = \llbracket S \rrbracket_p(\sigma)(\mathbf{x})$ (Lemma 14). The ϕ -expression allows us to pre-compute all the work that we may need to do when marginalising other discrete variables, which results in efficient nesting. Finally, the generation expression computes the conditional of a variable z given the rest of the parameters, as in Section 4.5 (see Lemma 15).

4.6.2 Eliminating a Single Variable z . The (ELIM GEN) rule below specifies a semantics-preserving transformation that takes a SlicStan program with a discrete MODEL-level parameter z , and transforms it to one where z is GENQUANT-level parameter. In practice, we apply this rule once per discrete MODEL-level parameter, which eliminates those parameters one-by-one, similarly to the variable elimination algorithm. And like in VE, the ordering in which we eliminate those variables can impact performance.

The (ELIM GEN) rule makes use of two auxiliary definitions that we define next. First, the neighbours of z , Γ_{ne} , are defined by the relation $\text{ne}(\Gamma, \Gamma', z)$ (Definition 7), which looks for non-data and non-continuous L1-variables in Γ' .

Definition 7 (Neighbours of z , $\text{ne}(\Gamma, \Gamma', z)$). For a \vdash typing environment Γ , a \vdash_2 typing environment $\Gamma' = \Gamma'_\sigma, \Gamma'_x$, and a variable $z \in \text{dom}(\Gamma'_x)$, the neighbours of z are defined as:

$$\text{ne}(\Gamma, \Gamma', z) \triangleq \{x : (\tau, \ell) \in \Gamma'_x \mid \ell = \mathbf{L1} \text{ and } \Gamma(x) = (\mathbf{int}\langle K \rangle, \mathbf{MODEL}) \text{ for some } K\}.$$

Second, $\text{st}(S_2)$ (Definition 8) is a statement that has the same store semantics as S_2 , but density semantics of 1: $\llbracket \text{st}(S_2) \rrbracket_s = \llbracket S_2 \rrbracket_s$, but $\llbracket \text{st}(S_2) \rrbracket_p = 1$. This ensures that the transformation preserves both the density semantics and the store semantics of S and is needed because $\mathbf{gen}(z)S_2$ discards any store computed by S_2 , thus only contributing to the weight.

Definition 8. Given a statement S , we define the statement $\text{st}(S)$ by replacing all $\mathbf{factor}(E)$ - and $L \sim d(E_1, \dots, E_n)$ -substatements in S by \mathbf{skip} (see Appendix A for the precise definition).

The elim-gen rule:

(ELIM GEN)

$$\begin{array}{c}
\Gamma(z) = (\mathbf{int}\langle K \rangle, \mathbf{MODEL}) \quad \Gamma_{\text{ne}} = \text{ne}(\Gamma, \Gamma_M, z) \quad S' = S_D; S'_M; S_Q \\
S'_M = S_1; f = \phi(\Gamma_{\text{ne}})\{\mathbf{elim}(\mathbf{int}\langle K \rangle z) S_2\}; \mathbf{factor}(f[\text{dom}(\Gamma_{\text{ne}})]); S_3; \mathbf{gen}(z)S_2; \text{st}(S_2) \\
\Gamma \vdash S : \mathbf{DATA} \quad S \Downarrow_{\Gamma} S_D, S_M, S_Q \quad \Gamma \xrightarrow{z} \Gamma_M \quad \Gamma_M \vdash_2 S_M : \mathbf{L1} \quad S_M \Downarrow_{\Gamma_M} S_1, S_2, S_3 \quad \Gamma' \vdash S' : \mathbf{DATA} \\
\hline
\Gamma, S \xrightarrow{z} \Gamma', S'
\end{array}$$

We can use the (ELIM GEN) rule to transform a SlicStan program, with respect to a parameter z , as described by Algorithm 1. This involves three main steps:

ALGORITHM 1: Single step of applying (ELIM GEN)

Arguments: $(\Gamma, S), z$	// A program (Γ, S) ; the variable z to eliminate
Requires: $\Gamma \vdash S : \text{DATA}$	// (Γ, S) is well-typed
Returns: (Γ', S')	// The transformed program

- 1: Slice (Γ, S) into S_D, S_M, S_Q according to $S \Downarrow_{\Gamma} S_D, S_M, S_Q$.
- 2: Derive incomplete Γ_M from Γ based on $\Gamma \xrightarrow{z} \Gamma_M$. // **DATA** of Γ is of level **L1** in Γ_M .
- 3: // Continuous **MODEL** var. of Γ are **L1** in Γ_M .
- 4: // z is of level **L2** in Γ_M .
- 5: // All other **MODEL** variables are given
- 6: // a type level placeholder in Γ_M .
- 7: Infer missing types of Γ_M according to $\Gamma_M \vdash_2 S_M : \text{L1}$.
- 8: Slice (Γ_M, S_M) into S_1, S_2, S_3 according to $S_M \Downarrow_{\Gamma_M} S_1, S_2, S_3$.
- 9:
- 10: $\Gamma_{ne} = ne(\Gamma, \Gamma_M, z)$ // Determine the discrete neighbours of z .
- 11: $S'_M = (S_1;$ // Eliminate z and re-generate z .
 - $f = \phi(\Gamma_{ne})\{\mathbf{elim}(\mathbf{int}\langle K \rangle z) S_2\};$
 - $\mathbf{factor}(f[\text{dom}(\Gamma_{ne})]);$
 - $S_3;$
 - $\mathbf{gen}(z) S_2;$
 - $\text{st}(S_2))$
- 12: $S' = S_D; S'_M; S_Q$
- 13: Infer an optimal Γ' according to $\Gamma' \vdash S' : \text{DATA}$
- 14: **return** (Γ', S')

- (1) Separate out S_M —the **MODEL**-level sub-part of S —using the top-level type system \vdash (line 1 of Algorithm 1).
- (2) Separate out S_2 —the part of S_M that involves the discrete parameter z —using the conditional independence type system \vdash_2 (lines 2–8).
- (3) Perform a single VE step by marginalising out z in S_2 and sample z from the conditional probability specified by S_2 (lines 10–11).

All other sub-statements of the program, $S_D, S_1, S_3,$ and $S_Q,$ stay the same during the transformation. By isolating S_2 and transforming only this part of the program, we make sure we do not introduce more work than necessary when performing variable elimination.

To efficiently marginalise out z , we want to find the *Markov boundary* of z given all **DATA** and continuous **MODEL** parameters: The data is given, and marginalisation happens inside the continuous parameters inference loop, so we can see continuous parameters as given for the purpose of discrete parameters marginalisation. Thus, we are looking for the relationship: $z \perp\!\!\!\perp \theta_{d2} \mid \mathcal{D}, \theta_c, \theta_{d1}$, where \mathcal{D} is the data, θ_c are the continuous **MODEL**-level parameters, θ_{d1} is a subset of the discrete **MODEL**-level parameters that is as small as possible (the Markov blanket), and θ_{d2} is the rest of the discrete **MODEL**-level parameters. We can find an optimal partitioning of the discrete parameters θ_{d1}, θ_{d2} that respects this relationship of interest using the type system from Section 3 together with type inference.

The judgement $\Gamma_M \vdash_2 S_M : \text{L1}$ induces a conditional independence relationship of the form $\mathbf{x}_{L2} \perp\!\!\!\perp \mathbf{x}_{L3} \mid \mathbf{x}_{L1}$, where $\mathbf{x} \models \Gamma_x$ (Theorem 3). The relation $\Gamma \xrightarrow{z} \Gamma_M$ (Definition 9) constrains the form of Γ_M based on Γ . This is needed to make sure we are working with a relationship of the form

we are interested in $-z \perp\!\!\!\perp \theta_{d2} \mid \mathcal{D}, \theta_c, \theta_{d1}$ —and that base types τ are the same between Γ and Γ_M . In particular, $\Gamma \xrightarrow{z} \Gamma_M$ constrains z to be the only **L2** parameter in Γ_M and all **DATA** and continuous **MODEL**-level parameters of Γ are **L1** in Γ_M . Note, $\text{dom}(\Gamma_M) \subseteq \text{dom}(\Gamma)$ and Γ_M only contains variables that are of level **MODEL** and below in Γ . Variables that are of level **GENQUANT** in Γ are not in Γ_M .

Definition 9 ($\Gamma \xrightarrow{z} \Gamma'$). For a \vdash_1 typing environment Γ and a \vdash_2 typing environment Γ' , a variable z and a statement S , we have:

$$\Gamma \xrightarrow{z} \Gamma' = \begin{cases} \Gamma(z) = (\tau, \text{MODEL}) \text{ and } \Gamma'_{x, \text{L2}} = \{z : \tau, \text{L2}\} \text{ for some } \tau \\ x : (\tau, \ell) \in \Gamma \text{ such that } \ell \leq \text{MODEL} \iff x : (\tau, \ell') \in \Gamma' \text{ for some } \ell' \in \{\text{L1}, \text{L2}, \text{L3}\} \\ x : (\tau, \text{DATA}) \in \Gamma \rightarrow x : (\tau, \text{L1}) \in \Gamma' \\ x : (\tau, \text{MODEL}) \in \Gamma_x \text{ and } \tau = \text{real} \text{ or } \tau = \text{real}[\dots] \rightarrow x : (\tau, \text{L1}) \in \Gamma'. \end{cases}$$

Following convention from earlier in the article, we use level subscripts to refer to specific subsets of Γ : In the above definition, $\Gamma'_{x, \text{L2}}$ refers to the subset of parameters x in Γ' , which are of level **L2**.

4.7 Marginalising Multiple Variables: An Example

To eliminate more than one discrete parameter, we apply the (**ELIM GEN**) rule repeatedly. Here, we work through a full example, showing the different steps of this repeated (**ELIM GEN**) transformation.

Consider an extended version of the HMM model from the beginning of this section (Program D), reformulated to include transformed parameters:

G. An extended HMM

```
S = real[2] phi ~ beta(1, 1);
real[2] theta ~ beta(1, 1);
real theta0 = theta[0];
int<2> z1 ~ bernoulli(theta0);
real theta1 = theta[z1];
int<2> z2 ~ bernoulli(theta1);
real theta2 = theta[z2];
int<2> z3 ~ bernoulli(theta2);
real phi1 = phi[z1];
real phi2 = phi[z2];
real phi3 = phi[z3];
data real y1 ~ normal(phi1, 1);
data real y2 ~ normal(phi2, 1);
data real y3 ~ normal(phi3, 1);
real theta3 = theta[z3];
int genz ~ bernoulli(theta4);
```

The typing environment

```
 $\Gamma = \{y_{1,2,3} : (\text{real}, \text{DATA}),$ 
 $\phi : (\text{real}[2], \text{MODEL}),$ 
 $\theta : (\text{real}[2], \text{MODEL}),$ 
 $\theta_{0,1,2} : (\text{real}, \text{MODEL}),$ 
 $\phi_{1,2,3} : (\text{real}, \text{MODEL}),$ 
 $z_{1,2,3} : (\text{int}<2>, \text{MODEL}),$ 
 $\theta_3 : (\text{real}, \text{GENQUANT}),$ 
 $\text{genz} : (\text{int}<2>, \text{GENQUANT})\}$ 
```

The variables we are interested in transforming are z_1, z_2 , and z_3 : These are the **MODEL**-level discrete parameters of Program G. The variable genz is already at **GENQUANT** level, so we can sample this with ancestral sampling (no need for automatic marginalisation).

We eliminate z_1, z_2 , and z_3 one-by-one, in that order. The order of elimination generally has a significant impact on the complexity of the resulting program (see also Section 4.8), but we do not focus on how to choose an ordering here. The problem of finding an optimal ordering is well-studied [Amir 2010; Arnborg et al. 1987; Kjærulff 1990] and is orthogonal to the focus of our work.

(1) Standard shredding of S

```

1   $S_D =$  skip;
2   $S_M =$  phi ~ beta(1, 1);
3      theta ~ beta(1, 1);
4      theta0 = theta[0];
5      z1 ~ bernoulli(theta0);
6      theta1 = theta[z1];
7      z2 ~ bernoulli(theta1);
8      theta2 = theta[z2];
9      z3 ~ bernoulli(theta2);
10     phi1 = phi[z1];
11     phi2 = phi[z2];
12     phi3 = phi[z3];
13     y1 ~ normal(phi1, 1);
14     y2 ~ normal(phi2, 1);
15     y3 ~ normal(phi3, 1);
16   $S_Q =$  theta3 = theta[z3];
17     genz ~ bernoulli(theta3);

```

(2) CI shredding of S_M

```

1   $S_1 =$  phi ~ beta(1, 1);
2      theta ~ beta(1, 1);
3      theta0 = theta[0];
4   $S_2 =$  z1 ~ bernoulli(theta0);
5      theta1 = theta[z1];
6      z2 ~ bernoulli(theta1);
7      phi1 = phi[z1];
8      y1 ~ normal(phi1, 1);
9   $S_3 =$  theta2 = theta[z2];
10     z3 ~ bernoulli(theta2);
11     phi2 = phi[z2];
12     phi3 = phi[z3];
13     y2 ~ normal(phi2, 1);
14     y3 ~ normal(phi3, 1);

```

(3) Applying (ELIM GEN): Program G-1

```

1  phi ~ beta(1, 1);
2  theta ~ beta(1, 1);
3  theta0 = theta[0];
4
5  f1 =  $\phi$ ([int<2> z2]){
6      elim(int<2> z1){
7          z1 ~ bernoulli(theta0);
8          theta1 = theta[z1];
9          z2 ~ bernoulli(theta1);
10         phi1 = phi[z1];
11         y1 ~ normal(phi1, 1);
12     }}
13  factor(f1[z2]);
14
15  theta2 = theta[z2];
16  z3 ~ bernoulli(theta2);
17  phi2 = phi[z2];
18  phi3 = phi[z3];
19  y2 ~ normal(phi2, 1);
20  y3 ~ normal(phi3, 1);
21
22  gen(int z1){
23      z1 ~ bernoulli(theta0);
24      theta1 = theta[z1];
25      z2 ~ bernoulli(theta1);
26      phi1 = phi[z1];
27      y1 ~ normal(phi1, 1);
28  }
29  theta1 = theta[z1];
30  phi1 = phi[z1];
31
32  theta3 = theta[z3];
33  genz ~ bernoulli(theta3);

```

Fig. 6. Step-by-step elimination of z_1 in Program G.

4.7.1 *Eliminating z_1* . Eliminating a single variable happens in three steps, as shown in Figure 6: standard shredding into S_D, S_M , and S_Q , conditional independence shredding of S_M into S_1, S_2 , and S_3 , and combining everything based on (ELIM GEN).

- (1) *Standard shredding*: $S \uparrow_{\Gamma} S_D, S_M, S_Q$. First, we separate out the parts of the program that depend on discrete parameters *generatively*; that is, any part of the program that would be in generated quantities with respect to the original program. In our case, this includes the last two lines in S . This would also include the **gen** parts of the transform program, which draw discrete parameters as generated quantities. Thus, $S \uparrow_{\Gamma} S_D, S_M, S_Q$, where S_D is empty, $S_Q = (\text{theta3} = \text{theta}[z3]; \text{genz} \sim \text{bernoulli}(\text{theta3}))$, and S_M is the rest of the program (see Figure 6, (1)).

- (2) *Conditional independence shredding*: $S_M \Downarrow_{\Gamma_M} S_1, S_2, S_3$. In the next step, we want to establish a conditional independence relationship $z_1 \perp\!\!\!\perp y, \phi_0, \theta_0, B$, where z_1 is some discrete parameter and A, B is a partitioning of the rest of the discrete parameters in the model: $\{z_2, z_3\}$. We derive a new, \vdash_2 typing environment Γ_M , using $\Gamma \xrightarrow{z} \Gamma_M$:

$$\begin{aligned} \Gamma_M = \{ & y_{1,2,3} : (\mathbf{real}, \mathbf{L1}), \phi : (\mathbf{real}[\mathbf{2}], \mathbf{L1}), \theta : (\mathbf{real}[\mathbf{2}], \mathbf{L1}), \\ & z_1 : (\mathbf{int}\langle\mathbf{2}\rangle, \mathbf{L2}), z_2 : (\mathbf{int}\langle\mathbf{2}\rangle, ?), z_3 : (\mathbf{int}\langle\mathbf{2}\rangle, ?), \\ & \theta_0 : (\mathbf{real}, \mathbf{L1}), \theta_{1,2} : (\mathbf{real}, ?), \phi_{1,2,3} : (\mathbf{real}, ?)\}. \end{aligned}$$

Here, we use the notation $?$ for a type placeholder, which will be inferred using type inference.

The optimal Γ_M under the type inference soft constraint $\mathbf{L3} < \mathbf{L1} < \mathbf{L2}$ such that $\Gamma_M \vdash_2 S_M : \mathbf{L1}$ is such that the levels of θ_1 and ϕ_1 are $\mathbf{L2}$, z_2 is $\mathbf{L1}$, and θ_2, ϕ_2 , and ϕ_3 are $\mathbf{L3}$. Shredding then gives us $S_M \Downarrow_{\Gamma_M} S_1, S_2, S_3$, as in Figure 6, (2).

- (3) *Combining based on (ELIM GEN)*. Having rearranged the program into suitable sub-statements, we use (ELIM GEN) to get Program G-1 (Figure 6, (3)) and:

$$\begin{aligned} \Gamma' = \{ & y_{1,2,3} : (\mathbf{real}, \mathbf{DATA}), \phi : (\mathbf{real}, \mathbf{MODEL}), \\ & \theta_1 : (\mathbf{real}, \mathbf{GENQUANT}), \theta_{0,2} : (\mathbf{real}, \mathbf{MODEL}), \\ & \phi_1 : (\mathbf{real}, \mathbf{GENQUANT}), \phi_{2,3} : (\mathbf{real}, \mathbf{MODEL}), \\ & z_1 : (\mathbf{int}, \mathbf{GENQUANT}), z_{2,3} : (\mathbf{int}\langle\mathbf{2}\rangle, \mathbf{MODEL}), \\ & \theta_3 : (\mathbf{real}, \mathbf{GENQUANT}), \text{genz} : (\mathbf{int}\langle\mathbf{2}\rangle, \mathbf{GENQUANT})\}. \end{aligned}$$

Eliminating z_2 . We apply the same procedure to eliminate the next variable, z_2 , from the updated Program G-1. The variable z_1 is no longer a \mathbf{MODEL} -level parameter, thus the only neighbouring parameter of z_2 is z_3 . Note also that the computation of the factor f_1 does not include any free discrete parameters (both z_1 and z_2 are local to the computation due to **elim** and ϕ). Thus, we do not need to include the computation of this factor anywhere else in the program (it does not get nested into other computations). We obtain a new program, Program G-2:

Program G-2	<pre> 18 }} 19 20 factor(f2[z3]); 21 phi3 = phi[z3]; 22 y3 ~ normal(phi3, 1); 23 24 gen(int z2){ 25 factor(f1[z2]); 26 theta2 = theta[z2]; 27 z3 ~ bernoulli(theta2); 28 phi2 = phi[z2]; 29 y2 ~ normal(phi2, 1); 30 } 31 theta2 = theta[z2]; 32 phi2 = phi[z2]; 33 34 gen(int z1){ 35 z1 ~ bernoulli(theta0); </pre>
<pre> 1 phi ~ beta(1, 1); 2 theta ~ beta(1, 1); 3 theta0 = theta[0]; 4 5 f1 = phi([int<2> z2]){ elim(int<2> z1){ 6 z1 ~ bernoulli(theta0); 7 theta1 = theta[z1]; 8 z2 ~ bernoulli(theta1); 9 phi1 = phi[z1]; 10 y1 ~ normal(phi1, 1); 11 }} 12 f2 = phi([int<2> z3]){ elim(int<2> z2){ 13 factor(f1[z2]); 14 theta2 = theta[z2]; 15 z3 ~ bernoulli(theta2); 16 phi2 = phi[z2]; 17 y2 ~ normal(phi2, 1); </pre>	

```

36  theta1 = theta[z1];
37  z2 ~ bernoulli(theta1);
38  phi1 = phi[z1];
39  y1 ~ normal(phi1, 1);
40  }
41  theta1 = theta[z1];
42  phi1 = phi[z1];
43
44  theta3 = theta0[z3];
45  genz ~ bernoulli(theta3);

```

Eliminating z_3 . Finally, we eliminate z_3 , which is the only discrete **MODEL**-level parameter left in the program. Thus, z_3 has no neighbours and f_3 is of arity 0: It is a real number instead of a vector. The final program generated by our implementation is Program G-3:

Program G-3

```

1  phi0 ~ beta(1, 1);
2  theta0 ~ beta(1, 1);
3
4  f1 =  $\phi(\text{int}\langle 2 \rangle z2)\{\text{elim}(\text{int}\langle 2 \rangle z1)\{$ 
5    z1 ~ bernoulli(theta0);
6    theta1 = theta[z1];
7    z2 ~ bernoulli(theta1);
8    phi1 = phi[z1];
9    y1 ~ normal(phi1, 1);
10 } }
11
12 f2 =  $\phi(\text{int}\langle 2 \rangle z3)\{\text{elim}(\text{int}\langle 2 \rangle z2)\{$ 
13   factor(f1[z2]);
14   theta2 = theta[z2];
15   z3 ~ bernoulli(theta2);
16   phi2 = phi[z2];
17   y2 ~ normal(phi2, 1);
18 } }
19
20 f3 =  $\phi()\{\text{elim}(\text{int}\langle 2 \rangle z3)\{$ 
21   factor(f2[z3]);
22   phi3 = phi[z3];
23   y3 ~ normal(phi3, 1);
24 } }
25
26 factor(f3);
27 gen(int z3){
28   factor(f2[z3]);
29   phi3 = phi[z3];
30   y3 ~ normal(phi3, 1);
31 }
32 phi3 = phi[z3];
33 gen(int z2){
34   factor(f1[z2]);
35   theta2 = theta[z2];
36   z3 ~ bernoulli(theta2);
37   phi2 = phi[z2];
38   y2 ~ normal(phi2, 1);
39 }
40 theta2 = theta[z2];
41 phi2 = phi[z2];
42 gen(int z1){
43   z1 ~ bernoulli(theta0);
44   theta1 = theta[z1];
45   z2 ~ bernoulli(theta1);
46   phi1 = phi[z1];
47   y1 ~ normal(phi1, 1);
48 }
49 theta1 = theta[z1];
50 phi1 = phi[z1];
51
52 gen3 = theta[z3];
53 genz ~ bernoulli(theta3);

```

4.8 Relating to Variable Elimination and Complexity Analysis

Assume \mathcal{D} , θ_d , and θ_c are the data, discrete model-level parameters, and continuous model-level parameters, respectively. As S_2 is a single-level statement of level **L2**, the density semantics of S_2 is of the form $\psi(\mathbf{x}_{L1}, \mathbf{x}_{L2}) = \psi(\mathcal{D}, \theta_c, \theta_{d,L1}, z)$ (Lemma 11).

As $\text{elim}(\text{int}\langle K \rangle z)$ binds the variable z and $\phi(\Gamma_{ne})\{_ \}$ binds the variables in $\text{dom}(\Gamma_{ne})$, the expression $\phi(\Gamma_{ne})\{\text{elim}(\text{int}\langle K \rangle z) S_2\}$ depends only on continuous parameters and data, and it contains no free mentions of any discrete variables. This means that the expression will be of level **L1** and shredded into S_1 during the marginalisation of any subsequent discrete variable z' . The substatement S_2 will always be some sub-statement of the original program (prior to any transformations), up to potentially several constant factors of the form $\text{factor}(f[\text{dom}(\Gamma_{ne})])$.

This observation makes it easy to reason about how repeated application of the (ELIM GEN) transform changes the complexity of the program. If the complexity of a SlicStan program with N discrete parameters of support $1, \dots, K$, is $O(S)$, then the complexity of a program where we naively marginalised out the discrete variables (Program E) will be $O(S \times K^N)$. In contrast, transforming with (ELIM GEN) gives us a program of complexity at most $O(N \times S \times K^{M+1})$ where M is the largest number of direct neighbours in the factor graph induced by the program. Further, the complexity could be smaller depending on the elimination ordering of choice. This result is not surprising, as we conjecture that repeated application of (ELIM GEN) is equivalent to variable elimination (though we do not formally prove this equivalence), which is of the same complexity.

It is clear from this complexity observation that VE is not always efficient. When the dependency graph is dense, M will be close to N , thus inference will be infeasible for large N . Fortunately, in many practical cases (such as those discussed in Section 5), this graph is sparse ($M \ll N$) and our approach is suitable and efficient. We note that this is a general limitation of exact inference of discrete parameters, and it is not a limitation of our approach in particular.

4.9 Semantic Preservation of the Discrete Variable Transformation

The result we are interested in is the semantic preservation of the transformation rule \xrightarrow{z} .

THEOREM 4 (SEMANTIC PRESERVATION OF \xrightarrow{z}). *For SlicStan programs Γ, S and Γ', S' , and a discrete parameter $z: \Gamma, S \xrightarrow{z} \Gamma', S'$ implies $\llbracket S \rrbracket = \llbracket S' \rrbracket$.*

PROOF. Note that shredding preserves semantics with respect to both \vdash and \vdash_2 (Lemma 6 and 10), examines the meaning of derived forms (Lemma 14 and 15), notes properties of single-level statements (Lemma 11), and applies the results on factorisation of shredding (Theorem 1) and conditional independence (Theorem 3). We present the full proof in Appendix A. \square

In addition, we also show that it is always possible to find a program derivable with (ELIM GEN), such that a MODEL-level variable z is transformed to a GENQUANT-level variable.

LEMMA 12 (EXISTENCE OF MODEL TO GENQUANT TRANSFORMATION). *For any SlicStan program Γ, S such that $\Gamma \vdash S : \mathbf{L1}$, and a variable $z \in \text{dom}(\Gamma)$ such that $\Gamma(z) = (\mathbf{int}\langle K \rangle, \mathbf{MODEL})$, there exists a SlicStan program Γ', S' , such that:*

$$\Gamma, S \xrightarrow{z} \Gamma', S' \quad \text{and} \quad \Gamma'(z) = (\mathbf{int}\langle K \rangle, \mathbf{GENQUANT}).$$

PROOF. By inspecting the level types of variables in each part of a program derivable using (ELIM GEN). We include the full proof in Appendix A. \square

The practical usefulness of Theorem 4 stems from the fact that it allows us to separate inference for discrete and continuous parameters. After applying (ELIM GEN) to each discrete MODEL-level parameter, we are left with a program that only has GENQUANT-level discrete parameters (Lemma 12). We can then slice the program into S_{HMC} and S_{GEN} and infer continuous parameters by using HMC (or other algorithms) on S_{HMC} and, next, draw the discrete parameters using ancestral sampling by running forward S_{GEN} . Theorem 4 tells us that this is a correct inference strategy.

When used in the context of a model with only discrete parameters, our approach corresponds to exact inference through VE. In the presence of discrete and continuous parameters, our transformation gives an analytical sub-solution for the discrete parameters in the model.

A limitation of our method is that, due to its density-based nature, it can only be applied to models of fixed size. It cannot, in its current form, support models where the number of random variables changes during inference, such as Dirichlet processes. However, this is a typical constraint adopted in Bayesian inference for efficiency. Another limitation is that discrete variables

Program H	Program H-A: Optimal transformation	Program H-B: Our transformation
<pre> data int K; real[K][K] phi; real[K] mu; int<K> z1 ~ categorical(phi[0]); int<K> z2; int<K> z3; if (z1 > K/2) { z2 ~ categorical(phi[z1]); z3 ~ categorical(phi[z2]); } else { z2 ~ categorical(phi[z1]); z3 ~ categorical(phi[z1]); } data real y1 ~ normal(mu[z1], 1); data real y2 ~ normal(mu[z2], 1); data real y3 ~ normal(mu[z3], 1); </pre>	<pre> ... factor(elim(int<2> z1) { z1 ~ categorical(phi[0]); y1 ~ normal(mu[z1], 1); if (z1 > K/2) { elim(int<2> z2) { elim(int<2> z3) { z2 ~ categorical(phi[z1]); z3 ~ categorical(phi[z2]); y2 ~ normal(mu[z2], 1); y3 ~ normal(mu[z3], 1); }} else { elim(int<2> z2) { z2 ~ categorical(phi[z1]); y2 ~ normal(mu[z2], 1); } elim(int<2> z3) { z3 ~ categorical(phi[z1]); y3 ~ normal(mu[z3], 1); }} }}); </pre>	<pre> ... f1 = ϕ(int<2> z2, int<2> z3){ elim(int<2> z1){ z1 ~ categorical(phi[0]); if(z1 > K/2){ z2 ~ categorical(phi[z1]); z3 ~ categorical(phi[z2]); } else{ z2 ~ categorical(phi[z1]); z3 ~ categorical(phi[z1]); } y1 ~ normal(mu[z1], 1); y2 ~ normal(mu[z2], 1); y3 ~ normal(mu[z3], 1); }} f2 = ϕ(int<2> z3){ elim(int<2> z2) factor(f1[z2, z3]); } f3 = ϕ(){ elim(int<2> z3) factor(f2[z3]); } factor(f3); </pre>

Fig. 7. A program with different conditional dependencies depending on control flow.

need to have finite (and fixed) support. For example, the method cannot be applied to transform a Poisson-distributed variable. In some but not all applications, truncating unbounded discrete parameters at a realistic upper bound would suffice to make the method applicable.

An advantage of our method is that it can be combined with any inference algorithm that requires a function proportional to the joint density of variables. This includes gradient-based algorithms, such as HMC and variational inference, but it can also be used with methods that allow for (e.g., unbounded) discrete variables as an analytical sub-solution that can optimise inference. For example, consider a Poisson variable $n \sim \text{Poisson}(\lambda)$ and a Binomial variable $k \sim \text{Binomial}(n, p)$. While n is of infinite support, and we cannot directly sum over all of its possible values, analytically marginalising out n gives us $k \sim \text{Poisson}(\lambda p)$. Future work can utilise such analytical results in place of explicit summation where possible.

4.10 Scope and Limitations of (ELIM GEN)

Previously, we discussed the scope of the conditional independence result of the article (Section 3.3). Similarly, here we demonstrate with an example, a situation where our approach of eliminating variables one-by-one using (ELIM GEN) is not optimal.

Consider the simple control-flow Program H below. In this example, z_2 and z_3 are *not* conditionally independent given $z_1 = 1$, but they are conditionally independent given $z_1 > K/2$. This independence is also referred to as context-specific independence [Boutilier et al. 1996; Minka and Winn 2009]. We can use different elimination strategy depending on which **if**-branch of the

program we find ourselves. Program H-A demonstrates this: Its complexity is $\mathcal{O}(\frac{K}{2} \times K^2 + \frac{K}{2} \times 2 \times K) = \mathcal{O}(\frac{1}{2}K^3 + K^2)$.

The typing relation \vdash_2 can only detect overall (in)dependencies, where sets of variables are conditionally independent given some X , regardless of what value X takes. Thus, our static analysis is not able to detect that $z_2 \perp\!\!\!\perp z_3 \mid z_1 = 0$. This results in Program H-B, which has complexity $\mathcal{O}(K^3 + K^2 + K)$: the same complexity as the optimal Program H-A, but with a bigger constant.

Even if we extend our approach to detect that z_2 and z_3 are independent in one branch, it is unclear how to incorporate this new information. Our strategy is based on computing intermediate factors that allow re-using already computed information: Eliminating z_1 requires computing a new factor f_1 that no longer depends on z_1 . We represented f_1 with a multidimensional array indexed by z_2 and z_3 , and we need to define each element of that array, thus we cannot decouple them for particular values of z_1 .

Runtime systems that compute intermediate factors in a similar way, such as Pyro [Uber AI Labs 2017], face that same limitation. Birch [Murray and Schön 2018], however, will be able to detect the conditional independence in the case $z_1 > K/2$, but it will not marginalise z_1 , as it cannot (analytically) marginalise over branches. Instead, it uses **Sequential Monte Carlo (SMC)** to repeatedly sample z_1 and proceed according to its value.

5 IMPLEMENTATION AND EMPIRICAL EVALUATION

The transformation we introduce can be useful for variety of models, and it can be adapted to PPLs to increase efficiency of inference and usability. Most notably, it can be used to extend Stan to allow for direct treatment of discrete variables, where previously that was not possible.

In this section, we present a brief overview of such a discrete parameter extension for SlicStan (Section 5.1). To evaluate the practicality of (ELIM GEN), we build a partial NumPyro [Phan et al. 2019] backend for SlicStan and compare our static approach to variable elimination for discrete parameters to the dynamic approach of NumPyro (Section 5.2). We find that our static transformation strategy speeds up inference compared to the dynamic approach, but that for models with a large number of discrete parameters performance gains could be diminished by the exponentially growing compilation time (Section 5.3).

In addition to demonstrating the practicality of our contribution through empirical evaluation, we also discuss the usefulness of our contribution through examples, in Appendix B.

5.1 Implementation

We update the original SlicStan¹⁴ according to the modification described in Section 2 and extend it to support automatic variable elimination through the scheme outlined in Section 4. As with the first version of SlicStan, the transformation produces a new SlicStan program that is then translated to Stan.

The variable elimination transformation procedure works by applying (ELIM GEN) iteratively, once for each discrete variable, as we show in Section 4.7. The level types `L1`, `L2`, and `L3` are not exposed to the user and are inferred automatically. Using bidirectional type-checking, we are able to synthesise a set of hard constraints that the levels must satisfy. These hard constraints will typically be satisfied by more than one assignment of variables to levels. We search for the optimal types with respect to the soft constraints `L3 < L1 < L2` using the theorem prover Z3 [De Moura and Bjørner 2008].

¹⁴Available at <https://github.com/mgorinova/SlicStan>.

5.2 Empirical Evaluation

To evaluate the practicality of our approach, we compare to the prior work most closely related to ours: that of Obermeyer et al. [2019], who implement efficient variable-elimination for plated factor graphs in Pyro [Uber AI Labs 2017]. Their approach uses effect-handlers and dynamically marginalises discrete variables, so gradient-based inference schemes can be used for the continuous parameters. This VE strategy has also been implemented in NumPyro [Phan et al. 2019].

As both ours and Pyro’s strategies correspond to VE, we do not expect to see differences in complexity of the resulting programs. However, as in our case the VE algorithm is determined and set up at compile time, while in the case of Pyro/NumPyro, this is done at runtime. The main question we aim to address is whether setting up the variable elimination logistics at compile time results in a practical runtime speed-up.

To allow for this comparison, we built a partial NumPyro backend for SlicStan. For each model we choose, we compare the runtime performance of three NumPyro programs:

- (1) The NumPyro program obtained by translating a SlicStan program with discrete parameters to NumPyro directly (labelled “NumPyro”). This is the baseline: We leave the discrete parameter elimination to NumPyro.
- (2) The NumPyro program obtained by translating a transformed SlicStan program, where all discrete parameters have been eliminated according to (ELIM GEN) (labelled “SlicStan”). The variable elimination setup is done at compile time; NumPyro does not do any marginalisation.
- (3) A hand-optimised NumPyro program, which uses the plate and markov program constructs to specify some of the conditional independencies in the program (labelled “NumPyro-Opt”).

In each case, we measure the time (in seconds) for sampling a single chain consisting of 2,500 warm-up samples and 10,000 samples using NUTS [Hoffman and Gelman 2014].

In addition, we report three compilation times:

- (1) The compilation time of the NumPyro program obtained by translating a SlicStan program with discrete parameters to NumPyro directly (labelled “NumPyro”).
- (2) The compilation time of the NumPyro program obtained by translating a transformed SlicStan program, where all discrete parameters have been eliminated (labelled “SlicStan”).
- (3) The time taken for the original SlicStan program to be transformed using (ELIM GEN) and translated to NumPyro code (labelled “SlicStan-to-NumPyro”).

We consider different numbers of discrete parameters for each model, up to 25 discrete parameters. We do not consider more than 25 parameters due to constraints of the NumPyro baseline, which we discuss in more detail in Section 5.3. We run experiments on two classes of model often seen in practice: hidden Markov models (Section 5.2.1) and mixture models (Section 5.2.2). To ensure a fair comparison, the same elimination ordering was used across experiments. Experiments were run on a dual-core 2.30 GHz Intel Xeon CPU and a Tesla T4 GPU (when applicable). All SlicStan models used in the experiments are available at the SlicStan repo.

5.2.1 Hidden Markov Models. We showed several examples of simple HMMs throughout the article (Program A, Program D, Program G) and worked through a complete example of VE in an HMM (4.7). We evaluate our approach on both the simple first-order HMM seen previously and on two additional ones: second-order HMM and factorial HMM.

First-order HMM. The first-order HMM is a simple chain of N discrete variables, each taking a value from 1 to K according to a categorical distribution. The event probabilities for the distribution of z_n are given by $\theta_{z_{n-1}}$, where θ is some given $K \times K$ matrix. Each data point \mathbf{y} is modelled




Fig. 8. HMM results.

as coming from a Gaussian distribution with mean μ_{z_n} and standard deviation 1, where μ is a K -dimensional continuous parameter of the model.

$$\begin{aligned} \mu_k &\sim \mathcal{N}(0, 1) \quad \text{for } k \in 1, \dots, K \\ z_1 &\sim \text{categorical}(\theta_1) \\ z_n &\sim \text{categorical}(\theta_{z_{n-1}}) \quad \text{for } n \in 2, \dots, N \\ y_n &\sim \mathcal{N}(\mu_{z_n}, 1) \quad \text{for } n \in 1, \dots, N \end{aligned}$$

We measure the compilation time and the time taken to sample 1 chain with each of the three NumPyro programs corresponding to this model. We use $K = 3$ and different values for N , ranging from $N = 3$ to $N = 25$. Figure 8 shows a summary of the results. We see that both on CPU and GPU, the program transformed using SlicStan outperforms the automatically generated NumPyro and also the manually optimised NumPyro-Opt. Each of the three programs has compilation time exponentially increasing with the number of variables, however SlicStan's compilation time increases the fastest. We discuss this drawback in more detail in Section 5.3, highlighting the importance of an extended loop-level analysis being considered in future work.

Second-order HMM. The second-order HMM is very similar to the first-order HMM, but the discrete variables depend on the previous two variables, in this case taking the maximum of the two.

$$\begin{aligned} \mu_k &\sim \mathcal{N}(0, 1) \quad \text{for } k \in 1, \dots, K \\ z_1 &\sim \text{categorical}(\theta_1), \quad z_2 \sim \text{categorical}(\theta_{z_1}) \\ z_n &\sim \text{categorical}(\theta_{\max(z_{n-2}, z_{n-1})}) \quad \text{for } n \in 3, \dots, N \\ y_n &\sim \mathcal{N}(\mu_{z_n}, 1) \quad \text{for } n \in 1, \dots, N \end{aligned}$$

Similarly to before, we run the experiment for $K = 3$ and different values for N , ranging from $N = 3$ to $N = 25$. We show the results in Figure 9, which once again shows SlicStan outperforming NumPyro and NumPyro-Opt in terms of runtime, but having slower compilation time for a larger number of discrete parameters.

Factorial HMM. In a factorial HMM, each data point y_n is generated using two independent hidden states z_n and h_n , each depending on the previous hidden states z_{n-1} and h_{n-1} .

$$\begin{aligned} \mu_k &\sim \mathcal{N}(0, 1) \quad \text{for } k \in 1, \dots, K^2 \\ z_1 &\sim \text{categorical}(\theta_1), \quad h_1 \sim \text{categorical}(\theta_1) \\ z_n &\sim \text{categorical}(\theta_{z_{n-1}}), \quad h_n \sim \text{categorical}(\theta_{h_{n-1}}) \quad \text{for } n \in 2, \dots, N \\ y_n &\sim \mathcal{N}(\mu_{z_n * h_n}, 1) \quad \text{for } n \in 1, \dots, N \end{aligned}$$




Fig. 9. Second-order HMM results.




Fig. 10. Factorial HMM results.

We run the experiment for $K = 3$ and different length of the chain N , ranging from $N = 1$ (two discrete parameters) to $N = 12$ (24 discrete parameters). We show the results in Figure 10: Similarly to before, SlicStan outperforms both NumPyro and NumPyro-Opt in terms of runtime. We also observe that, in the case of SlicStan, the time taken to sample a single chain increases more slowly as we increase the number of discrete variables.

5.2.2 Mixture Models. Another useful application of mixed discrete and continuous variable models is found in mixture models. We run experiments on two models: soft K -means clustering and linear regression with outlier detection.

Soft K -means. The Gaussian mixture model underlines the celebrated soft K -means algorithm. Here, we are interested in modelling some D -dimensional data that belongs to one of K (unknown) Gaussian clusters. Each cluster k is specified by a D -dimensional mean $\mu_{.,k}$. Each data point $y_{.,n}$ is associated with a cluster z_n .

$$\begin{aligned} \mu_{d,k} &\sim \mathcal{N}(0, 1) \quad \text{for } d \in 1, \dots, D \text{ and } k \in 1, \dots, K \\ z_n &\sim \text{categorical}(\pi) \quad \text{for } n \in 1, \dots, N \\ y_{d,n} &\sim \mathcal{N}(\mu_{d,z_n}, 1) \quad \text{for } d \in 1, \dots, D \text{ and } n \in 1, \dots, N \end{aligned}$$

We run the experiments for $K = 3, D = 10$, and $N = 3, \dots, 25$ and show the results in Figure 11.

We observe a clear linear trend of the runtime growing with N , with SlicStan performing better and its runtime growing more slowly than that of NumPyro. While the SlicStan-translated code runs faster than NumPyro-Opt for $N \leq 25$, we observe that the SlicStan runtime grows faster than that of the manually optimised NumPyro-Opt code.

Outlier detection. The final model we consider is a Bayesian linear regression that allows for outlier detection. The model considers data points (x_n, y_n) , where y lies on the line $\alpha x + \beta$ with some added noise. The noise σ_{z_n} depends on a Bernoulli parameter z_n , which corresponds to whether or not the point (x_n, y_n) is an outlier or not. The noise for outliers (σ_1) and the noise for non-outliers




Fig. 11. Soft K-means results.




Fig. 12. Outliers results.

(σ_2) are given as hyperparameters.

$$\alpha \sim \mathcal{N}(0, 10), \quad \beta \sim \mathcal{N}(0, 10)$$

$$\pi_1^{(raw)} \sim \mathcal{N}(0, 1), \quad \pi_2^{(raw)} \sim \mathcal{N}(0, 1), \quad \pi = \frac{\exp \pi_1^{(raw)}}{\exp \pi_1^{(raw)} + \exp \pi_2^{(raw)}}$$

$$z_n \sim \text{bernoulli}(\pi) \quad \text{for } n \in 1, \dots, N$$

$$y_n \sim \mathcal{N}(\alpha x_n + \beta, \sigma_{z_n}) \quad \text{for } n \in 1, \dots, N$$

Similarly to the earlier HMM models, SlicStan has the smallest runtime per chain, but at the expense of fast growing compile time (Figure 12).

5.3 Analysis and Discussion

Our method can be applied to general models containing a fixed and known number of finite-support discrete parameters, which significantly reduces the amount of manual effort that was previously required for such models in languages like Stan [Damiano et al. 2018]. In addition, as shown in Figures 8–12, SlicStan outperforms both the NumPyro baseline and the hand-optimised NumPyro-Opt in terms of runtime. This suggests that a static-time discrete variable optimisation, like the one introduced in this article, is indeed beneficial and speeds up inference.

One limitation of our experimental analysis is the relatively small number of discrete parameters we consider. Due to the array dimension limit imposed by PyTorch/NumPy, Pyro cannot have more than 25 discrete variables (64 for CPU) unless the dependence between them is specified using markov or plate (as with NumPyro-Opt). For NumPyro this hardcoded limit is 32. Thus, it would not be possible to compare to the NumPyro baseline for a larger number of variables, though comparing to the hand-optimised NumPyro-Opt would still be possible.

Perhaps the biggest limitation of the discrete parameters version of SlicStan is the exponentially growing compilation time. Using a semi-lattice instead of a lattice in the \vdash_2 level type analysis breaks the requirement of the bidirectional type system that ensures efficiency of type

inference. The constraints generated by the type system can no longer be resolved by SlicStan’s original linear-time algorithm. While polynomial-time constraint-solving strategy may still exist, we choose to employ Z3 to automatically resolve the type inference constraints and leave the consideration for more efficient type inference algorithm for future work.

This also highlights the importance of a future SlicStan version that considers *arrays* of discrete parameters. Our algorithm currently supports only individual discrete parameters. In the cases where the size of an array of discrete parameters is statically known, the (ELIM GEN) procedure can be applied to a program where such arrays have been “flattened” into a collection of individual discrete variables, which is the strategy we adopt for the experiments in this section. But to be applicable more widely, the (ELIM GEN) rule needs to be generalised based on array element level dependence analysis, for example, by incorporating ideas from the polyhedral model [Fautrier 1992]. As the array level dependence analysis that would be required in most practical use-cases is very straightforward, we believe this would be a useful and feasible applied extension of our work. In addition, this would significantly decrease the number of program variables for which we need to infer a level type during the (ELIM GEN) transformation, thus making compilation practical for larger or arbitrary numbers of discrete parameters.

6 RELATED WORK

This article provides a type system that induces conditional independence relationships, and it discusses one practical application of such type system: an automatic marginalisation procedure for discrete parameters of finite support.

Conditional independence. The theoretical aim of our article is similar to that of Barthe et al. [2019], who discuss a separation logic for reasoning about independence, and the follow-up work of Bao et al. [2021], who extend the logic to capture conditional independence. One advantage of our method is that the verification of conditional independence is automated by type inference, while it would rely on manual reasoning in the works of Barthe et al. [2019] and Bao et al. [2021]. However, the logic approach can be applied to a wider variety of verification tasks. Amtoft and Banerjee [2020] show a correspondence between variable independence and slicing a discrete-variables-only probabilistic program. The biggest difference to our work is that their work considers only conditional independence of variables given the observed data: that is, CI relationships of the form $\mathbf{x}_1 \perp\!\!\!\perp \mathbf{x}_2 \mid \mathcal{D}$ for some subsets of variables \mathbf{x}_1 and \mathbf{x}_2 and data \mathcal{D} . The language of Amtoft and Banerjee [2020] requires observed data to be specified *syntactically* using an **observe** statement. Conditional independencies are determined only given this observed data, and the method for determining how to slice a program is tied to the **observe** statements. From the Amtoft and Banerjee [2020] paper: “A basic intuition behind our approach is that an **observe** statement can be removed if it does not depend on something on which the returned variable x also depends.” In contrast, we are able to find CI relationships given any variables we are interested in ($\mathbf{x}_1 \perp\!\!\!\perp \mathbf{x}_2 \mid \mathbf{x}_3$ for some \mathbf{x}_1 , \mathbf{x}_2 , and \mathbf{x}_3), and type inference constitutes of a straightforward algorithm for finding such relationships. However, Amtoft and Banerjee [2020] permit unbounded number of variables (e.g., **while** ($y > 0$) $y \sim \text{bernoulli}(0.2)$), while it is not clear how to extend SlicStan/Stan to support this. While not in a probabilist programming setting, Lobo-Vesga et al. [2020] use taint analysis to find independencies between variables in a program to facilitate easy trade off between privacy and accuracy in differential privacy context.

Automatic marginalisation. The most closely related previous work, in terms of the automatic marginalisation procedure, is that of Obermeyer et al. [2019] and that of Murray et al. [2018]. Obermeyer et al. [2019] implement efficient variable-elimination for plated factor graphs in Pyro [Uber AI Labs 2017]. Their approach uses effect-handlers and can be implemented in other

effect-handling-based PPLs, such as Edward2 [Tran et al. 2018]. Murray et al. [2018] introduce a “delayed sampling” procedure in Birch [Murray and Schön 2018], which optimises the program via partial analytical solutions to sub-programs. Their method corresponds to automatic variable elimination and, more generally, automatic Rao–Blackwellization. While we focus on discrete variable elimination only, our conditional independence type system can be directly used for more general analysis. The method from Section 4 can be extended to marginalise out and sample continuous variables whenever they are part of an analytically tractable sub-program, similarly to delayed sampling in Birch. One key difference of our approach is that the program re-writes are guided by the type system and happen at compile time, before inference is run. In contrast, both Pyro and Birch maintain a dynamic graph that guides the analysis at runtime.

Symbolic inference. Where a full analytical solution is possible, several probabilistic programming languages can derive it via symbolic manipulation, including Hakaru [Narayanan et al. 2016] and PSI [Gehr et al. 2016, 2020], while Dice [Holtzen et al. 2020] performs exact inference for models with discrete parameters only, by analysing the program structure. In contrast, we focus on re-writing the program and decomposing it into parts to be used with fast and more general asymptotically exact or approximate inference algorithms, such as HMC, variational inference, or others.

Extending HMC to support discrete parameters. The idea of modifying HMC to handle discrete variables and discontinuities has been previously explored [Nishimura et al. 2017; Pakman and Paninski 2013; Zhang et al. 2012; Zhou 2020]. More recently, Zhou et al. [2019] introduced the probabilistic programming language LF-PPL, which is designed specifically to be used with the **Discontinuous Hamiltonian Monte Carlo (DHMC)** algorithm [Nishimura et al. 2017]. The algorithm and their framework can also be extended to support discrete parameters. LF-PPL provides support for an HMC version that itself works with discontinuities. Our approach is to statically rewrite the program to match the constraints of Stan, vanilla HMC and its several well-optimised extensions, such as NUTS [Hoffman and Gelman 2014].

Composable and programmable inference. Recent years have seen a growing number of techniques that allow for tailored-to-the-program compilation to an inference algorithm. For example, Gen [Cusumano-Towner et al. 2019] can statically analyse the model structure to compile to a more efficient inference strategy. In addition, languages such as Gen and Turing [Ge et al. 2018] facilitate composable and programmable inference [Mansinghka et al. 2018], where the user is provided with inference building blocks to implement their own model-specific algorithm. Our method can be understood as an automatic composition between two inference algorithms: variable elimination and HMC or any other inference algorithm that can be used to sample continuous variables.

7 CONCLUSION

This article introduces an information flow type system that can be used to check and infer conditional independence relationships in probabilistic programs, through type checking and inference, respectively. We present a practical application of this type system: a semantics-preserving transformation that makes it possible to use, and to efficiently and automatically infer discrete parameters in SlicStan, Stan, and other density-based probabilistic programming languages. The transformed program can be seen as a hybrid inference algorithm on the original program, where continuous parameters can be drawn using efficient gradient-based inference methods, like HMC, while the discrete parameters are drawn using variable elimination.

While the variable elimination transformation uses results on conditional independence of discrete parameters, our type system is not restricted to this usage. Conditional independence relationships can be of interest in many contexts in probabilistic modelling, including more general use of variable elimination, message-passing algorithms, Rao-Blackwellization, and factorising a program for a composed-inference approach. We believe conditional independence by typing can enable interesting future work that automates the implementation of such methods.

APPENDICES

A DEFINITIONS AND PROOFS

A.1 Definitions

Definition 10 (Assigns-to Set $W(S)$). $W(S)$ is the set that contains the names of global variables that have been assigned to within the statement S . It is defined recursively as follows:

$$\begin{aligned} W(x[E_1] \dots [E_n] = E) &= \{x\} & W(\mathbf{skip}) &= \emptyset \\ W(S_1; S_2) &= W(S_1) \cup W(S_2) & W(\mathbf{factor}(E)) &= \emptyset \\ W(\mathbf{if}(E) S_1 \mathbf{else} S_2) &= W(S_1) \cup W(S_2) & W(L \sim d(E_1, \dots, E_n)) &= \emptyset. \\ W(\mathbf{for}(x \mathbf{in} E_1 : E_2) S) &= W(S) \setminus \{x\} \end{aligned}$$

Definition 11 (Reads Set $R(S)$). $R(S)$ is the set that contains the names of global variables that have been read within the statement S . It is defined recursively as follows:

$$\begin{aligned} R(x) &= \{x\} & R(x[E_1] \dots [E_n] = E) &= \bigcup_{i=1}^n R(E_i) \cup R(E) \\ R(c) &= \emptyset & R(S_1; S_2) &= R(S_1) \cup R(S_2) \\ R([E_1, \dots, E_n]) &= \bigcup_{i=1}^n R(E_i) & R(\mathbf{if}(E) S_1 \mathbf{else} S_2) &= R(E) \cup R(S_1) \cup R(S_2) \\ R(E_1[E_2]) &= R(E_1) \cup R(E_2) & R(\mathbf{for}(x \mathbf{in} E_1 : E_2) S) &= R(E_1) \cup R(E_2) \cup R(S) \setminus \{x\} \\ R(f(E_1, \dots, E_n)) &= \bigcup_{i=1}^n R(E_i) & R(\mathbf{skip}) &= \emptyset \\ R([E|x \mathbf{in} E_1 : E_2]) &= R(E) \cup R(E_1) \cup R(E_2) & R(\mathbf{factor}(E)) &= R(E) \\ R(\mathbf{target}(S)) &= R(S) & R(L \sim d(E_1, \dots, E_n)) &= R(L) \cup R(E_1) \cup \dots \cup R(E_n). \end{aligned}$$

Definition 12 (Samples-to Set $\widetilde{W}(S)$). $\widetilde{W}(S)$ is the set that contains the names of global variables that have been sampled within the statement S . It is defined recursively as follows:

$$\begin{aligned} \widetilde{W}(L = E) &= \emptyset & \widetilde{W}(\mathbf{skip}) &= \emptyset \\ \widetilde{W}(S_1; S_2) &= \widetilde{W}(S_1) \cup \widetilde{W}(S_2) & \widetilde{W}(\mathbf{factor}(E)) &= \emptyset \\ \widetilde{W}(\mathbf{if}(E) S_1 \mathbf{else} S_2) &= \widetilde{W}(S_1) \cup \widetilde{W}(S_2) & \widetilde{W}(x[E_1] \dots [E_n] \sim d(E_1, \dots, E_n)) &= \{x\}. \\ \widetilde{W}(\mathbf{for}(x \mathbf{in} E_1 : E_2) S) &= \widetilde{W}(S) \setminus \{x\} \end{aligned}$$

Definition 13 (Free Variables $FV(S)$). $FV(S)$ is the set that contains the free variables that are used in a statement S . It is recursively defined as follows:

$$\begin{aligned} FV(x) &= \{x\} & FV(S_1; S_2) &= FV(S_1) \cup FV(S_2) \\ FV(c) &= \emptyset & FV(\mathbf{if}(E) S_1 \mathbf{else} S_2) &= FV(E) \cup FV(S_1) \cup FV(S_2) \\ FV([E_1, \dots, E_n]) &= \bigcup_{i=1}^n FV(E_i) & FV(\mathbf{for}(x \mathbf{in} E_1 : E_2) S) &= FV(E_1) \cup FV(E_2) \cup \\ FV(E_1[E_2]) &= FV(E_1) \cup FV(E_2) & FV(S) \setminus \{x\} & \\ FV(f(E_1, \dots, E_n)) &= \bigcup_{i=1}^n FV(E_i) & FV(\mathbf{skip}) &= \emptyset \\ FV([E|x \mathbf{in} E_1 : E_2]) &= FV(E) \cup FV(E_1) \cup FV(E_2) & FV(\mathbf{factor}(E)) &= FV(E) \\ FV(\mathbf{target}(S)) &= FV(S) & FV(L \sim d(E_1, \dots, E_n)) &= FV(L) \cup FV(E_1) \cup \dots \cup \\ FV(x[E_1] \dots [E_n] = E) &= \bigcup_{i=1}^n FV(E_i) \cup FV(E) & FV(E_n). & \end{aligned}$$

Definition 14. We overload the notation $\Gamma(L)$ that looks up the type of an L-value in Γ . When applied to a more general expression E , $\Gamma(E)$ looks up the *type level* of E in Γ :

$$\begin{aligned} \Gamma(x) &= \ell, \text{ where } \ell \text{ is the level of } x \text{ in } \Gamma & \Gamma(E_1[E_2]) &= \Gamma(E_1) \sqcup \Gamma(E_2) \\ \Gamma(c) &= \mathbf{DATA} & \Gamma(f(E_1, \dots, E_n)) &= \bigsqcup_{i=1}^n \Gamma(E_i) \\ \Gamma([E_1, \dots, E_n]) &= \bigsqcup_{i=1}^n \Gamma(E_i) & \Gamma([E]x \text{ in } E_1 : E_2) &= \Gamma(E) \sqcup \Gamma(E_1) \sqcup \Gamma(E_2). \end{aligned}$$

Definition 15. $\Gamma(E_1, \dots, E_n) \equiv \Gamma(E_1) \sqcup \dots \sqcup \Gamma(E_n)$.

Definition 16 ($R_{\Gamma, \ell}(S)$). $R_{\Gamma, \ell}(S)$ is the set that contains the names of global variables that have been read at level ℓ within the statement S . It is defined recursively as follows:

$$\begin{aligned} R_{\Gamma, \ell}(x[E_1] \dots [E_n] = E) &= \begin{cases} \bigcup_{i=1}^n R(E_i) \cup R(E) & \text{if } \Gamma(x) = (_, \ell) \\ \emptyset & \text{otherwise,} \end{cases} \\ R_{\Gamma, \ell}(S_1; S_2) &= R_{\Gamma, \ell}(S_1) \cup R_{\Gamma, \ell}(S_2) \\ R_{\Gamma, \ell}(\mathbf{if}(E) S_1 \mathbf{else} S_2) &= R_{\Gamma, \ell}(E) \cup R_{\Gamma, \ell}(S_1) \cup R_{\Gamma, \ell}(S_2) \\ R_{\Gamma, \ell}(\mathbf{for}(x \text{ in } E_1 : E_2) S) &= R_{\Gamma, \ell}(E_1) \cup R_{\Gamma, \ell}(E_2) \cup R_{\Gamma, \ell}(S) \setminus \{x\} \\ R_{\Gamma, \ell}(\mathbf{skip}) &= \emptyset \\ R_{\Gamma, \ell}(\mathbf{factor}(E)) &= \begin{cases} R(E) & \text{if } \ell = \mathbf{MODEL} \\ \emptyset & \text{else} \end{cases} \\ R_{\Gamma, \ell}(L \sim d(E_1, \dots, E_n)) &= \begin{cases} R(L \sim d(E_1, \dots, E_n)) & \text{if } \ell = \bigsqcup \{\ell' \mid \exists x \in FV(L \sim d(E_1, \dots, E_n)) \\ \exists \tau. \Gamma(x) = (\tau, \ell')\} \\ \emptyset & \text{otherwise.} \end{cases} \end{aligned}$$

Definition 17 ($W_{\Gamma, \ell}(S)$). $W_{\Gamma, \ell}(S) \triangleq \{x \in W(S) \mid \Gamma(x) = (\tau, \ell) \text{ for some } \tau\}$.

Definition 18 ($\widetilde{W}_{\Gamma, \ell}(S)$). $\widetilde{W}_{\Gamma, \ell}(S) \triangleq \{x \in \widetilde{W}(S) \mid \Gamma(x) = (\tau, \ell) \text{ for some } \tau\}$.

Definition 19. Given a statement S , we define the statement $\text{st}(S)$ by structural induction on S :

$$\begin{aligned} \text{st}(x[E_1] \dots [E_n] = E) &= x[E_1] \dots [E_n] = E \\ \text{st}(S_1; S_2) &= \text{st}(S_1); \text{st}(S_2) \\ \text{st}(\mathbf{if}(E) S_1 \mathbf{else} S_2) &= \mathbf{if}(E) \text{st}(S_1) \mathbf{else} \text{st}(S_2) \\ \text{st}(\mathbf{for}(x \text{ in } E_1 : E_2) S) &= \mathbf{for}(x \text{ in } E_1 : E_2) \text{st}(S) \\ \text{st}(\mathbf{skip}) &= \mathbf{skip} \\ \text{st}(\mathbf{factor}(E)) &= \mathbf{skip} \\ \text{st}(L \sim d(E_1, \dots, E_n)) &= \mathbf{skip}. \end{aligned}$$

Definition 20 (*Neighbours of z , $\text{ne}(\Gamma, \Gamma', z)$*). For a \vdash typing environment Γ , a \vdash_2 typing environment $\Gamma' = \Gamma'_\sigma, \Gamma'_x$ and a variable $z \in \text{dom}(\Gamma'_x)$, the neighbours of z are defined as:

$$\text{ne}(\Gamma, \Gamma', z) \triangleq \{x : (\tau, \ell) \in \Gamma'_x \mid \ell = \mathbf{L1} \text{ and } \Gamma(x) = (\mathbf{int}\langle K \rangle, \mathbf{MODEL}) \text{ for some } K\}.$$

A.2 Proofs

RESTATEMENT OF LEMMA 1 (NONINTERFERENCE OF \vdash). *Suppose $s_1 \models \Gamma$, $s_2 \models \Gamma$, and $s_1 \approx_\ell s_2$ for some ℓ . Then for *SlicStan* statement S and expression E :*

- (1) *If $\Gamma \vdash E : (\tau, \ell)$ and $(s_1, E) \Downarrow V_1$ and $(s_2, E) \Downarrow V_2$, then $V_1 = V_2$.*
- (2) *If $\Gamma \vdash S : \ell$ and $(s_1, S) \Downarrow s'_1, w_1$ and $(s_2, S) \Downarrow s'_2, w_2$, then $s'_1 \approx_\ell s'_2$.*

PROOF. (1) follows by rule induction on the derivation $\Gamma \vdash E : (\tau, \ell)$, and using that if $\Gamma \vdash E : (\tau, \ell)$, $x \in R(E)$ and $\Gamma(x) = (\tau', \ell')$, then $\ell' \leq \ell$. (2) follows by rule induction on the derivation $\Gamma \vdash S : \ell$ and using (1).

Most cases follow trivially from the inductive hypothesis. An exception is the **(TARGET)** case, which we show below.

We use the premise $\forall \ell' > \ell. R_{\Gamma+\ell'}(S) = \emptyset$, together with a lemma that for S, s_1 and s_2 such that $s_1, S \Downarrow s'_1, w_1$, and $s_2, S \Downarrow s'_2, w_2$, and $\forall x \in R(S). s_1(x) = s_2(x)$, we have that $w_1 = w_2$. (This lemma follows by structural induction on S .) In the case of **(TARGET)**, $s_1, \text{target}(S) \Downarrow w_1$, and $s_2, \text{target}(S) \Downarrow w_2$ and $R(S) = \bigcup_{\ell'} R_{\Gamma+\ell'}(S) = (\bigcup_{\ell' \leq \ell} R_{\Gamma+\ell'}(S)) \cup (\bigcup_{\ell' > \ell} R_{\Gamma+\ell'}(S)) = \bigcup_{\ell' \leq \ell} R_{\Gamma+\ell'}(S)$. Then, for any $x \in R(S)$, $x \in R_{\Gamma+\ell'}(S)$ for some $\ell' \leq \ell$, so $\Gamma(x) = (\tau, \ell_x)$ such that $\ell_x \leq \ell' \leq \ell$. And thus, by definition of \approx_ℓ , $s_1(x) = s_2(x)$ for any $x \in R(S)$. By applying the lemma above, we then get $w_1 = w_2$, as required. \square

RESTATEMENT OF LEMMA 4 (SHREDDING PRODUCES SINGLE-LEVEL STATEMENTS).

$$S \Downarrow_{\Gamma} (S_D, S_M, S_Q) \implies \Gamma \vdash \text{DATA}(S_D) \wedge \Gamma \vdash \text{MODEL}(S_M) \wedge \Gamma \vdash \text{GENQUANT}(S_Q)$$

PROOF. By rule induction on the derivation of $S \Downarrow_{\Gamma} S_D, S_M, S_Q$. \square

RESTATEMENT OF LEMMA 5 (PROPERTY OF SINGLE-LEVEL STATEMENTS). *Let $\Gamma_\sigma, \Gamma_x \vdash S$ be SlicStan program, such that S is single-level statement of level ℓ , $\Gamma \vdash \ell(S)$. Then there exist unique functions f and ϕ , such that for any $\sigma, \mathbf{x} \models \Gamma_\sigma, \Gamma_x$:*

$$\llbracket S \rrbracket(\sigma)(\mathbf{x}) = f(\sigma_{\leq \ell}, \mathbf{x}_{\leq \ell}) \cup \sigma_{> \ell}, \quad \phi(\sigma_{\leq \ell})(\mathbf{x}_{\leq \ell}),$$

where we write $\sigma_{\leq \ell} = \{(x \mapsto V) \in \sigma \mid \Gamma_\sigma(x) = (_, \ell)\}$ and $\sigma_{> \ell} = \sigma \setminus \sigma_{\leq \ell}$.

PROOF. This property follows from noninterference (Lemma 1), if we understand factor and sample statements as assignments to a reserved weight variables of different levels. Let Γ, S be a SlicStan program and suppose we obtain S' by:

- Substituting every **factor**(E) statement with $w_\ell = w_\ell * E$, where $\Gamma(E) = \text{real}, \ell$ and w_{DATA} , w_{MODEL} , and w_{GENQUANT} are write-only, distinct, and reserved variables in the program.
- Substituting every $L \sim d(E_1, \dots, E_n)$ statement with $w_\ell = w_\ell * d_{\text{pdf}}(L \mid E_1, \dots, E_n)$, where $\Gamma(d_{\text{pdf}}(L \mid E_1, \dots, E_n)) = \text{real}, \ell$.

Then for all $\sigma, \mathbf{x} \models \Gamma$, we have $\llbracket S \rrbracket_p(\sigma)(\mathbf{x}) = \prod_{\ell} \sigma'(w_\ell)$, where $\sigma' = \llbracket S' \rrbracket_s(\sigma, \forall \ell. w_\ell \mapsto 1)(\mathbf{x})$. By non-interference (Lemma 1), for any level ℓ and store $\sigma_2 \approx_\ell \sigma$, if $\sigma'_2 = \llbracket S' \rrbracket_s(\sigma_2, \forall \ell. w_\ell \mapsto 1)(\mathbf{x})$, then $\sigma'_2 \approx_\ell \sigma'$. Thus, $\sigma'_2(w_{\ell'}) = \sigma_2(w_{\ell'})$ for $\ell' \leq \ell$, and therefore, when S is a single-level statement of level ℓ , $\llbracket S \rrbracket_s(\sigma, \forall \ell. w_\ell \mapsto 1)(\mathbf{x}) = f(\sigma_{\leq \ell}, \mathbf{x}_{\leq \ell}), \sigma_{> \ell}, w_{\leq \ell} \mapsto \phi(\sigma_{\leq \ell}, \mathbf{x}_{\leq \ell}), w_{> \ell} \mapsto 1$, for some functions f and ϕ . Finally, this gives us $\llbracket S \rrbracket_s(\sigma, \mathbf{x}) = (f(\sigma_{\leq \ell}, \mathbf{x}_{\leq \ell}), \sigma_{> \ell})$, $\llbracket S \rrbracket_p(\sigma, \mathbf{x}) = \phi(\sigma_{\leq \ell}, \mathbf{x}_{\leq \ell})$. \square

RESTATEMENT OF LEMMA 6 (SEMANTIC PRESERVATION OF \Downarrow_{Γ}). *If $\Gamma \vdash S : \text{DATA}$ and $S \Downarrow_{\Gamma} (S_D, S_M, S_Q)$ then $\llbracket S \rrbracket = \llbracket S_D; S_M; S_Q \rrbracket$.*

PROOF. Follows by adapting proof from Gorinova et al. [2019]. \square

RESTATEMENT OF LEMMA 10 (SEMANTIC PRESERVATION OF $\Downarrow_{\Gamma} 2$). *If $\Gamma \vdash_2 S : \text{L1}$ and $S \Downarrow_{\Gamma} S_1, S_2, S_3$, then $\llbracket S \rrbracket = \llbracket S_1; S_2; S_3 \rrbracket$.*

PROOF. Follows by adapting proof from Gorinova et al. [2019]. \square

LEMMA 13. *For a SlicStan expression E and a function $\phi(x, y) = V$, where V is a value such that $(\sigma, x, y), E \Downarrow V$ for every x and y and some σ , if $x \notin R(E)$, then:*

$$\exists \phi' \text{ such that } \phi(x, y) = \phi'(y) \text{ for all } x, y.$$

PROOF. By induction on the structure of E . □

RESTATEMENT OF THEOREM 1 (SHREDDING INDUCES A FACTORISATION OF THE DENSITY). *Suppose $\Gamma \vdash S : \text{DATA}$ and $S \Downarrow_{\Gamma} S_D, S_M, S_Q$ and $\Gamma = \Gamma_{\sigma} \cup \Gamma_{\mathcal{D}} \cup \Gamma_{\theta} \cup \Gamma_Q$. For all $\sigma, \mathcal{D}, \theta$, and Q : if $\sigma, \mathcal{D}, \theta, Q \models \Gamma_{\sigma}, \Gamma_{\mathcal{D}}, \Gamma_{\theta}, \Gamma_Q$, and $\llbracket S \rrbracket_p(\sigma)(\mathcal{D}, \theta, Q) \propto p(\mathcal{D}, \theta, Q)$ and $\widetilde{W}(S_Q) = \text{dom}(\Gamma_Q)$, then:*

- (1) $\llbracket S_M \rrbracket_p(\sigma_D)(\mathcal{D}, \theta, Q) \propto p(\theta, \mathcal{D})$,
- (2) $\llbracket S_Q \rrbracket_p(\sigma_M)(\mathcal{D}, \theta, Q) = p(Q \mid \theta, \mathcal{D})$,

where $\sigma_D = \llbracket S_D \rrbracket_s(\sigma)(\mathcal{D}, \theta, Q)$ and $\sigma_M = \llbracket S_M \rrbracket_s(\sigma_D)(\mathcal{D}, \theta, Q)$.

PROOF. We prove this by establishing a more general result:

For $\sigma, \mathcal{D}, \theta, Q \models \Gamma_{\sigma}, \Gamma_{\mathcal{D}}, \Gamma_{\theta}, \Gamma_Q$, $A = \widetilde{W}(S_Q) \subseteq Q$ and some $B \subseteq Q \setminus A$, if $\llbracket S \rrbracket_p(\sigma)(\mathcal{D}, \theta, Q) \propto p(\mathcal{D}, \theta, A \mid B)$, then:

- (1) $\llbracket S_D \rrbracket_p(\sigma)(\mathcal{D}, \theta, Q) = 1$,
- (2) $\llbracket S_M \rrbracket_p(\sigma_D)(\mathcal{D}, \theta, Q) = p(\theta, \mathcal{D})$,
- (3) $\llbracket S_Q \rrbracket_p(\sigma_M)(\mathcal{D}, \theta, Q) = p(A \mid \theta, \mathcal{D}, B)$.

Note that in the case where $\widetilde{W}(S_Q) = Q$, we have $A = Q$ and $B = \emptyset$, and the original statement of the theorem, $\llbracket S_Q \rrbracket_p(\sigma_M)(\mathcal{D}, \theta, Q) = p(Q \mid \theta, \mathcal{D})$, holds.

We prove the extended formulation above by induction on the structure of S and use of Lemma 2, Lemma 4 and Lemma 5, Lemma 6.

Take any $\sigma, \mathcal{D}, \theta, Q \models \Gamma_{\sigma}, \Gamma_{\mathcal{D}}, \Gamma_{\theta}, \Gamma_Q$ and let

$$\begin{aligned} \Phi(S, S_D, S_M, S_Q) &\triangleq \\ \Gamma \vdash S : \text{data} \wedge S \Downarrow_{\Gamma} S_D, S_M, S_Q \wedge A = \widetilde{W}(S_Q) & \\ \implies \exists B \subseteq Q \setminus A. \forall \sigma_D, \sigma_M. (& \\ \llbracket S \rrbracket_p(\sigma)(\mathcal{D}, \theta, Q) \propto p(\mathcal{D}, \theta, A \mid B) \wedge \llbracket S_D \rrbracket(\sigma)(\mathcal{D}, \theta, Q) = \sigma_D \wedge \llbracket S_M \rrbracket(\sigma_D)(\mathcal{D}, \theta, Q) = \sigma_M & \\ \implies \llbracket S_D \rrbracket_p(\sigma)(\mathcal{D}) = 1 & \\ \wedge \llbracket S_M \rrbracket_p(\sigma_D)(\mathcal{D}, \theta) = p(\theta, \mathcal{D}) & \\ \wedge \exists B \subseteq Q \setminus \widetilde{W}(S_Q). \llbracket S_Q \rrbracket_p(\sigma_M)(\mathcal{D}, \theta, Q) = p(A \mid \theta, \mathcal{D}, B) & \end{aligned}$$

Take any Γ, S, S_D, S_M, S_Q such that $S \Downarrow_{\Gamma} S_D, S_M, S_Q$, $A = \widetilde{W}(S_Q)$, and take any $\sigma, \mathcal{D}, \theta, Q \models \Gamma_{\sigma}, \Gamma_{\mathcal{D}}, \Gamma_{\theta}, \Gamma_Q$, an unnormalised density p and $B \subseteq Q \setminus A$, such that $\llbracket S \rrbracket_p(\sigma)(\mathcal{D}, \theta, Q) \propto p(\mathcal{D}, \theta, A \mid B)$. We prove by rule induction on the derivation of $S \Downarrow_{\Gamma} S_D, S_M, S_Q$ that $\Phi(S, S_D, S_M, S_Q)$.

(SHRED SEQ). Let $S = S_1; S_2$ and $S_1 \Downarrow_{\Gamma} S_{1D}, S_{1M}, S_{1Q}$ and $S_2 \Downarrow_{\Gamma} S_{2D}, S_{2M}, S_{2Q}$. Thus, $S \Downarrow_{\Gamma} (S_{1D}; S_{2D}), (S_{1M}; S_{2M}), (S_{1Q}; S_{2Q})$.

Assume $\Phi(S_1, S_{1D}, S_{1M}, S_{1Q})$ and $\Phi(S_2, S_{2D}, S_{2M}, S_{2Q})$.

Let:

- $A_1 = \widetilde{W}(S_{1Q})$ and $B_1 \subseteq Q \setminus A_1$ is such that $\llbracket S_{1Q} \rrbracket_p(\sigma_M)(\mathcal{D}, \theta, Q) = p_1(A_1 \mid \mathcal{D}, \theta, B_1)$.
- $\llbracket S_1 \rrbracket(\sigma)(\mathcal{D}, \theta, Q) = \sigma'$.
- $\llbracket S_1 \rrbracket_p(\sigma)(\mathcal{D}, \theta, Q) \propto p_1(\mathcal{D}, \theta, A_1 \mid B_1)$.
- $A_2 = \widetilde{W}(S_{2Q})$ and $B_2 \subseteq Q \setminus A_2$ is such that $\llbracket S_{2Q} \rrbracket_p(\sigma_M)(\mathcal{D}, \theta, Q) = p_2(A_2 \mid \mathcal{D}, \theta, B_2)$.
- $\llbracket S_2 \rrbracket_p(\sigma')(\mathcal{D}, \theta, Q) \propto p_2(\mathcal{D}, \theta, A_2 \mid B_2)$.

Thus, by Lemma 2, $\llbracket S \rrbracket_p = \llbracket S_1; S_2 \rrbracket_p = \llbracket S_1 \rrbracket_p \times \llbracket S_2 \rrbracket_p$, so $p(\mathcal{D}, \theta, A \mid B) \propto p_1(\mathcal{D}, \theta, A_1 \mid B_1) p_2(\mathcal{D}, \theta, A_2 \mid B_2)$.

For (1), we have $\forall \sigma \models \Gamma_{\sigma}. \llbracket S_{1D} \rrbracket_p(\sigma)(\mathcal{D}, \theta, Q) = \llbracket S_{2D} \rrbracket_p(\sigma)(\mathcal{D}, \theta, Q) = 1$. Thus, by Lemma 2, $\llbracket S_{1D}; S_{2D} \rrbracket_p = \llbracket S_{1D} \rrbracket_p \times \llbracket S_{2D} \rrbracket_p = 1$.

From $\Phi(S_1, S_{1D}, S_{1M}, S_{1Q})$ and $\Phi(S_2, S_{2D}, S_{2M}, S_{2Q})$, we also have:

- $\llbracket S_{1Q} \rrbracket_p(\sigma_M)(\mathcal{D}, \theta, Q) = p(A_1 \mid \theta, \mathcal{D}, B_1)$,
- $\llbracket S_{2Q} \rrbracket_p(\sigma'_M)(\mathcal{D}, \theta, Q) = p(A_2 \mid \theta, \mathcal{D}, B_2)$,

$$A = \widetilde{W}(S_Q) = \widetilde{W}(S_{1Q}; S_{2Q}) = \widetilde{W}(S_{1Q}) \cup \widetilde{W}(S_{2Q}) = A_1 \cup A_2.$$

From S well typed, it must be the case that $A_1 \cap A_2 = \emptyset$. Thus, we write $A = A_1, A_2$.

We will prove that the property holds for $B = B_1 \cup B_2 \setminus A_1 \setminus A_2$.

By semantic preservation of \Updownarrow_Γ (Lemma 6), $\llbracket S_1 \rrbracket_p = \llbracket S_{1D}; S_{1M}; S_{1Q} \rrbracket_p = \llbracket S_{1D} \rrbracket_p \times \llbracket S_{1M} \rrbracket_p \times \llbracket S_{1Q} \rrbracket_p \propto 1 \times p_1(\theta, \mathcal{D}) \times p_1(A_1 \mid \theta, \mathcal{D}, B_1)$. Similarly, $\llbracket S_2 \rrbracket_p \propto 1 \times p_2(\theta, \mathcal{D}) \times p_2(A_2 \mid \theta, \mathcal{D}, B_2) = p_2(\theta, \mathcal{D})p_2(A_2 \mid \theta, \mathcal{D}, A_1, B_1)$.

But $p(\theta, \mathcal{D}, A \mid B) \propto p_1(\theta, \mathcal{D}, A_1 \mid B_1)p_2(\theta, \mathcal{D}, A_2 \mid B_2)$, so:

$$p(\theta, \mathcal{D}, A \mid B) \propto p_1(\theta, \mathcal{D})p_1(A_1 \mid \theta, \mathcal{D}, B_1)p_2(\theta, \mathcal{D})p_2(A_2 \mid \theta, \mathcal{D}, A_1, B_1).$$

So,

$$\begin{aligned} p(\theta, \mathcal{D}) &= \int p(\mathcal{D}, \theta, A \mid B)p(B)dAdB \\ &\propto \int p_1(\theta, \mathcal{D})p_1(A_1 \mid \theta, \mathcal{D}, B_1)p_2(\theta, \mathcal{D})p_2(A_2 \mid \theta, \mathcal{D}, A_1, B_1)p(B)dA_1dA_2dB \\ &\propto p_1(\theta, \mathcal{D})p_2(\theta, \mathcal{D}) \int p(B)p_1(A_1 \mid \theta, \mathcal{D}, B_1)p_2(A_2 \mid \theta, \mathcal{D}, A_1, B_1)dA_1dA_2dB \\ &= p_1(\theta, \mathcal{D})p_2(\theta, \mathcal{D}) \int p(B) \left(\int p_1(A_1 \mid \theta, \mathcal{D}, B_1) \left(\int p_2(A_2 \mid \theta, \mathcal{D}, A_1, B_1)dA_2 \right) dA_1 \right) dB \\ &= p_1(\theta, \mathcal{D})p_2(\theta, \mathcal{D}) \\ &\propto p_1(\theta, \mathcal{D})p_2(\theta, \mathcal{D}). \end{aligned}$$

$$\text{Thus, } \llbracket S_M \rrbracket_p = \llbracket S_{1M}; S_{2M} \rrbracket_p \propto p_1(\theta, \mathcal{D})p_2(\theta, \mathcal{D}) \propto p(\theta, \mathcal{D}).$$

Finally, for last property on S , we use the chain rule of probability, semantics property of sequencing, and the result from above to get:

$$\begin{aligned} p(A \mid \mathcal{D}, \theta, B) &= \frac{p(\mathcal{D}, \theta, A \mid B)}{p(\mathcal{D}, \theta \mid B)} \\ &\propto \frac{p_1(\mathcal{D}, \theta)p_2(\mathcal{D}, \theta)p_1(A_1 \mid \mathcal{D}, \theta, B_1)p_2(A_2 \mid \mathcal{D}, \theta, B_2)}{p(\mathcal{D}, \theta)} \times \frac{p(B)}{p(B \mid \mathcal{D}, \theta)} \\ &\propto p_1(A_1 \mid \mathcal{D}, \theta, B_1)p_2(A_2 \mid \mathcal{D}, \theta, B_2) \\ &= \llbracket S_{1Q} \rrbracket_p \llbracket S_{2Q} \rrbracket_p = \llbracket S_Q \rrbracket_p. \end{aligned}$$

Thus:

$$p(A \mid \mathcal{D}, \theta, B) = \frac{p_1(A_1 \mid \mathcal{D}, \theta, B_1)p_2(A_2 \mid \mathcal{D}, \theta, B_2)}{Z},$$

where:

$$\begin{aligned} Z &= \int p_1(A_1 \mid \mathcal{D}, \theta, B_1)p_2(A_2 \mid \mathcal{D}, \theta, B_2)dA \\ &= \int p_1(A_1 \mid \mathcal{D}, \theta, B_1) \left(\int p_2(A_2 \mid \mathcal{D}, \theta, B_2)dA_2 \right) dA_1 \\ &= 1. \end{aligned}$$

So, $Z = 1$, and $p(A \mid \mathcal{D}, \theta, B) = p_1(A_1 \mid \mathcal{D}, \theta, B_1)p_2(A_2 \mid \mathcal{D}, \theta, B_2) = \llbracket S_Q \rrbracket_p$.

Thus:

- $\llbracket S_D \rrbracket_p = \llbracket S_{1D}; S_{2D} \rrbracket_p = 1,$
- $\llbracket S_M \rrbracket_p = \llbracket S_{1M}; S_{2M} \rrbracket_p \propto p_1(\theta, \mathcal{D})p_2(\theta, \mathcal{D}) = p(\theta, \mathcal{D}),$
- $\llbracket S_Q \rrbracket_p = \llbracket S_{1Q}; S_{2Q} \rrbracket_p = p_1(A_1 \mid \theta, \mathcal{D}, B_1)p_2(A_2 \mid \theta, \mathcal{D}, A_1, B_1) = p(A_1, A_2 \mid \theta, \mathcal{D}, B),$

$\Phi((S_1; S_2), (S_{1D}; S_{2D}), (S_{1M}; S_{2M}), (S_{1Q}; S_{2Q}))$ from here. \square

RESTATEMENT OF LEMMA 9 (SHREDDING PRODUCES SINGLE-LEVEL STATEMENTS 2).

$$S \Downarrow_{\Gamma} S_1, S_2, S_3 \implies \Gamma \vdash \mathbf{L1}(S_1) \wedge \Gamma \vdash \mathbf{L2}(S_2) \wedge \Gamma \vdash \mathbf{L3}(S_3).$$

PROOF. By rule induction on the derivation of $S \Downarrow_{\Gamma} S_1, S_2, S_3$. \square

RESTATEMENT OF LEMMA 10 (SEMANTIC PRESERVATION OF \Downarrow_{Γ} , 2). *If $\Gamma \vdash_2 S : \mathbf{L1}$ and $S \Downarrow_{\Gamma} S_1, S_2, S_3$, then $\llbracket S \rrbracket = \llbracket S_1; S_2; S_3 \rrbracket$.*

PROOF. \square

RESTATEMENT OF LEMMA 11 (PROPERTY OF SINGLE-LEVEL STATEMENTS 2). *Let $\Gamma_{\sigma}, \Gamma_{\mathbf{x}}, S$ be a SlicStan program, and $\Gamma \vdash_2 S : \mathbf{L1}$, and S is single-level statement of level ℓ , $\Gamma \vdash_2 \ell(S)$. Then there exist unique functions f and ϕ , such that for any $\sigma, \mathbf{x} \models \Gamma_{\sigma}, \Gamma_{\mathbf{x}}$:*

- (1) *If $\ell = \mathbf{L1}$, then $\llbracket S \rrbracket(\sigma)(\mathbf{x}) = (f(\sigma_{\mathbf{L1}}, \mathbf{x}_{\mathbf{L1}}), \sigma_{\mathbf{L2}}, \sigma_{\mathbf{L3}}), \phi(\sigma_{\mathbf{L1}})(\mathbf{x}_{\mathbf{L1}}),$*
- (2) *If $\ell = \mathbf{L2}$, then $\llbracket S \rrbracket(\sigma)(\mathbf{x}) = (\sigma_{\mathbf{L1}}, f(\sigma_{\mathbf{L1}}, \sigma_{\mathbf{L2}}, \mathbf{x}_{\mathbf{L1}}, \mathbf{x}_{\mathbf{L2}}), \sigma_{\mathbf{L3}}), \phi(\sigma_{\mathbf{L1}}, \sigma_{\mathbf{L2}})(\mathbf{x}_{\mathbf{L1}}, \mathbf{x}_{\mathbf{L2}}),$*
- (3) *If $\ell = \mathbf{L3}$, then $\llbracket S \rrbracket(\sigma)(\mathbf{x}) = (\sigma_{\mathbf{L1}}, \sigma_{\mathbf{L2}}, f(\sigma_{\mathbf{L1}}, \sigma_{\mathbf{L3}}, \mathbf{x}_{\mathbf{L1}}, \mathbf{x}_{\mathbf{L3}})), \phi(\sigma_{\mathbf{L1}}, \sigma_{\mathbf{L3}})(\mathbf{x}_{\mathbf{L1}}, \mathbf{x}_{\mathbf{L3}}).$*

PROOF. By understanding factor and sample statements as assignment to a reserved weight variables of different levels (similarly to Lemma 5) and noninterference (Lemma 7). \square

RESTATEMENT OF LEMMA 12 (EXISTENCE OF MODEL TO GENQUANT TRANSFORMATION). *For any SlicStan program Γ, S such that $\Gamma \vdash S : \mathbf{L1}$, and a variable $z \in \text{dom}(\Gamma)$ such that $\Gamma(z) = (\mathbf{int}\langle K \rangle, \mathbf{MODEL})$, there exists a SlicStan program Γ', S' , such that,*

$$\Gamma, S \xrightarrow{z} \Gamma', S' \quad \text{and} \quad \Gamma'(z) = (\mathbf{int}\langle K \rangle, \mathbf{GENQUANT}).$$

PROOF. Take a SlicStan program Γ, S , a typing environment Γ_M , a variable z , and statements S_D, S_M , and S_Q , such that:

$$\Gamma(z) = (\mathbf{int}\langle K \rangle, \mathbf{MODEL}) \quad \Gamma \vdash S : \mathbf{DATA} \quad \Gamma \xrightarrow{z} \Gamma_M \quad S \Downarrow_{\Gamma} S_D, S_M, S_Q \quad \Gamma_M \vdash_2 S_M : \mathbf{L1}.$$

Take also statements S_1, S_2, S_3 , and S'_M , and a typing environment Γ_{ne} such that

$$S_M \Downarrow_{\Gamma_M} S_1, S_2, S_3 \quad \Gamma_{\text{ne}} = \text{ne}(\Gamma, \Gamma_M, z),$$

$$S'_M = S_1; f = \phi(\Gamma_{\text{ne}})\{\mathbf{elim}(\mathbf{int}\langle K \rangle z) S_2\}; \mathbf{factor}(f[\text{dom}(\Gamma_{\text{ne}})]); S_3; \mathbf{gen}(z)S_2; \text{st}(S_2).$$

Let Γ' is such that $\text{dom}(\Gamma') = \text{dom}(\Gamma) \cup \{f\}$ and for all $x : \tau, \ell \in \Gamma$:

$$\Gamma'(x) = \begin{cases} (\tau, \ell) & \text{if } \ell \neq \mathbf{MODEL}, \\ (\tau, \ell) & \text{if } \ell = \mathbf{MODEL} \text{ and } \Gamma_M(x) \neq (\tau, \mathbf{L2}), \\ (\tau, \mathbf{GENQUANT}) & \text{if } \ell = \mathbf{MODEL} \text{ and } \Gamma_M(x) = (\tau, \mathbf{L2}). \end{cases}$$

By semantic preservation of shredding (Lemma 6, Lemma 10) and type preservation of the operational semantics (Gorinova et al. [2019]), $\Gamma \vdash S_D; S_1; S_2; S_3; S_Q : \mathbf{DATA}$, and thus, by (SEQ), $\Gamma \vdash S_D : \mathbf{DATA}, \Gamma \vdash S_1 : \mathbf{DATA}, \dots, \Gamma \vdash S_Q : \mathbf{DATA}$.

By definition of $\Gamma', \Gamma'_{\mathbf{DATA}} \subset \Gamma_{\mathbf{DATA}}$. S_D is single-level of level \mathbf{DATA} and $\Gamma \vdash S_D : \mathbf{DATA}$, so $\Gamma_{\mathbf{DATA}} \vdash S_D : \mathbf{DATA}$ and thus $\Gamma' \vdash S_D : \mathbf{DATA}$. Similarly, $\Gamma \vdash S_1 : \mathcal{D}$ and $\Gamma \vdash S_3 : \mathcal{D}$.

$\Gamma \vdash S_2 : \mathbf{DATA}$, so using (PHI), (ELIM), and (FACTOR), and noting that by definition $\text{dom}(\Gamma_{\text{ne}}) \subset \text{dom}(\Gamma_{M, \mathbf{L1}})$, so $\Gamma_{\text{ne}} \subset \Gamma$, we can derive:

$$\Gamma' \vdash f = \phi(\Gamma_{\text{ne}})\{\mathbf{elim}(\mathbf{int}\langle K \rangle z) S_2\}; \mathbf{factor}(f[\text{dom}(\Gamma_{\text{ne}})]) : \mathbf{DATA}.$$

By $\Gamma \vdash S_2 : \mathbf{DATA}$ and the definition of Γ' , and using (GEN) and definition of st , we also derive:

$$\Gamma' \vdash \mathbf{gen}(z) S_2; \text{st}(S_2) : \mathbf{GENQUANT}.$$

Finally, S_Q is a single-level statement of level $\mathbf{GENQUANT}$ and for all $x : \tau, \ell \in \Gamma, x : \tau, \ell' \in \Gamma$, where $\ell \leq \ell'$. Therefore, $\Gamma \vdash S_Q : \mathbf{DATA}$ implies $\Gamma' \vdash S_Q : \mathbf{DATA}$.

Altogether, this gives us $\Gamma' \vdash S_D; S'_M; S_Q$, and so by (ELIM GEN), $\Gamma, S \xrightarrow{z} \Gamma', S_D; S'_M, S_Q$. \square

LEMMA 14. Let Γ, S be a SlicStan program, such that $\sigma, \mathbf{x} \models \Gamma, \llbracket S \rrbracket_s(\sigma)(\mathbf{x}) = \sigma'$ and $\llbracket S \rrbracket_p(\sigma)(\mathbf{x}) = \psi(\mathbf{x})$ for some function ψ . If $f \notin \text{dom}(\Gamma)$ is a fresh variable, then $z, z_1, \dots, z_n \in \text{dom}(\Gamma_x)$ are discrete variables of base types $\mathbf{int}\langle K \rangle, \mathbf{int}\langle K_1 \rangle, \dots, \mathbf{int}\langle K_n \rangle$, respectively, and S' is a statement such that

$$S' = f = \phi(\mathbf{int}\langle K_1 \rangle z_1, \dots, \mathbf{int}\langle K_n \rangle z_n)\{\mathbf{elim}(\mathbf{int}\langle K \rangle z) S\}; \mathbf{factor}(f[z_1, \dots, z_n]);$$

then $\llbracket S' \rrbracket_s(\sigma)(\mathbf{x}) = \sigma''$ with $\sigma''[-f] = \sigma'$ and $\llbracket S' \rrbracket_p(\sigma)(\mathbf{x}) = \sum_{z=1}^K \psi(\mathbf{x})$.

PROOF. By examining the operational semantics of assignment, \mathbf{factor} , and the derived forms \mathbf{elim} and ϕ . \square

LEMMA 15. Let Γ, S be a SlicStan program, such that $\sigma, \mathbf{x} \models \Gamma, \llbracket S \rrbracket_s(\sigma)(\mathbf{x}) = \sigma'$ and $\llbracket S \rrbracket_p(\sigma)(\mathbf{x}) = \psi(\mathbf{x})$ for some function ψ . If $z \in \text{dom}(\Gamma_x)$ is a discrete variable of base type $\mathbf{int}\langle K \rangle$, and S' is a statement such that

$$S' = \mathbf{gen}(z) S; \text{st}(S);$$

then $\llbracket S' \rrbracket_s(\sigma)(\mathbf{x}) = \sigma', \psi(\mathbf{x})$ is normalisable with respect to z with $\psi(\mathbf{x}) \propto p(z \mid \mathbf{x} \setminus \{z\})$, and $\llbracket S' \rrbracket_p(\sigma)(\mathbf{x}) = p(z \mid \mathbf{x} \setminus \{z\})$.

PROOF. By examining the operational semantics of \sim and \mathbf{target} , and by induction on the structure of S to prove $\llbracket \text{st}(S) \rrbracket_s = \llbracket S \rrbracket_s$ and $\llbracket \text{st}(S) \rrbracket_p = 1$. \square

Typing Rules for Derived Forms:

(ELIM)

$$\frac{\Gamma' \vdash S : \mathbf{DATA} \quad R_{\Gamma, \mathbf{GENQUANT}}(S) = \emptyset \quad \Gamma' = \Gamma[z \mapsto \mathbf{int}\langle K \rangle, \mathbf{MODEL}]}{\Gamma \vdash \mathbf{elim}(\mathbf{int}\langle K \rangle z) S : \mathbf{MODEL}}$$

(GEN)

$$\frac{\Gamma(z) = (\mathbf{int}, \mathbf{GENQUANT}) \quad \Gamma \vdash S : \mathbf{DATA}}{\Gamma \vdash \mathbf{gen}(\mathbf{int}\langle K \rangle z) S : \mathbf{GENQUANT}}$$

(PHI)

$$\frac{\Gamma' \vdash S : \mathbf{DATA} \quad \forall \ell' > \ell. R_{\Gamma \vdash \ell'}(S) = \emptyset \quad \Gamma' = \Gamma[z_1 \mapsto (\mathbf{int}\langle K_1 \rangle, \ell), \dots, z_N \mapsto (\mathbf{int}\langle K_N \rangle, \ell)]}{\Gamma \vdash \phi(\mathbf{int}\langle K_1 \rangle z_1, \dots, \mathbf{int}\langle K_N \rangle z_N) S : \mathbf{real}, \ell}$$

RESTATEMENT OF THEOREM 4 (SEMANTIC PRESERVATION OF \xrightarrow{z}). For SlicStan programs Γ, S and Γ', S' , and a discrete parameter $z: \Gamma, S \xrightarrow{z} \Gamma', S' \rightarrow \llbracket S \rrbracket = \llbracket S' \rrbracket$.

PROOF. Let Γ, S and Γ', S' be SlicStan programs, and z be a discrete parameter, such that $\Gamma, S \xrightarrow{z} \Gamma', S'$. Let $S \Downarrow_{\Gamma} S_D, S_M, S_Q, S \Downarrow_{\Gamma'} S'_D, S'_M, S'_Q$, and $S_M \Downarrow_{\Gamma''} S_1, S_2, S_3$ for Γ'' such that $\Gamma \xrightarrow{z} \Gamma''$ and $\Gamma'' \vdash_2 S_M : \mathbf{L1}$.

Let $\Gamma = \Gamma_{\sigma}, \Gamma_{\mathbf{DATA}}, \Gamma_{\mathbf{MODEL}}, \Gamma_{\mathbf{GENQUANT}}, \Gamma' = \Gamma'_{\sigma}, \Gamma'_{\mathbf{DATA}}, \Gamma'_{\mathbf{MODEL}}, \Gamma'_{\mathbf{GENQUANT}}$ and $\Gamma'' = \Gamma''_{\sigma}, \Gamma''_{\mathbf{L1}}, \Gamma''_{\mathbf{L2}}, \Gamma''_{\mathbf{L3}}$ be the usual partitioning of each of the typing environments.

Let z be a store such that $z \models \{z : \Gamma(z)\}$.

Let \mathcal{D}, θ and Q be stores such that $\mathcal{D} \models \Gamma_{\mathbf{DATA}}, z, \theta \models \Gamma_{\mathbf{MODEL}}$, and $Q \models \Gamma_{\mathbf{GENQUANT}}$.

Let θ_1, θ_2 and θ_3 be a partitioning of θ , such that $\mathcal{D}, \theta_1 \models \Gamma''_{\mathbf{L1}}, z, \theta_2 \models \Gamma''_{\mathbf{L2}}$, and $\theta_3 \models \Gamma''_{\mathbf{L3}}$.

Then, by definition of $\Gamma \xrightarrow{z} \Gamma'', \theta_2 = z$.

By Theorem 1:

- $\llbracket S_D \rrbracket_p(\sigma)(\mathcal{D}, z, \theta, Q) = 1$,
- $\llbracket S_M \rrbracket_p(\sigma_D)(\mathcal{D}, z, \theta, Q) \propto p(z, \theta, \mathcal{D})$,
- $\llbracket S_Q \rrbracket_p(\sigma_M)(\mathcal{D}, z, \theta, Q) = p(Q \mid z, \theta, \mathcal{D})$.

$\Gamma, S \xrightarrow{d} \Gamma', S'$, thus S' must be of the form

$$S' = S_D; S_1; f = \phi(\Gamma_{\mathbf{L1}}'')\{\mathbf{elim}(\mathbf{int}(K)z) S_2\}; \mathbf{factor}(f[\mathbf{dom}(\Gamma_{\mathbf{L1}}'')]); S_3; \mathbf{gen}(z)S_2; \mathbf{st}(S_2); S_Q,$$

where $\Gamma \vdash S : \mathbf{DATA}, S \Downarrow_{\Gamma} S_D, S_M, S_Q, \Gamma \xrightarrow{z} \Gamma'', \Gamma \vdash_2 S_M : \mathbf{L1}$, and $S_M \Downarrow_{\Gamma''} S_1, S_2, S_3$.

The relation \Downarrow_{Γ} is semantics-preserving for well-typed programs with respect to both \vdash and \vdash_2 (Lemma 6 and Lemma 10). Thus, $\llbracket S \rrbracket = \llbracket S_D; S_1; S_2; S_3; S_Q \rrbracket$.

We present a diagrammatic derivation of the change on store and density that each sub-part in the original and transformed program makes in Figure 13.

Combining all of these results gives that:


$$\llbracket S' \rrbracket_s(\sigma)(\mathcal{D}, \theta, Q) = \sigma'' = \sigma'[f \mapsto v] = \llbracket S \rrbracket_s(\sigma)((\mathcal{D}, \theta, Q))[f \mapsto v].$$

In other words, the transformation \xrightarrow{z} preserves store semantics (up to creating of one new fresh variable f).

For the density, we get:

$$\begin{aligned} & \llbracket S' \rrbracket_p(\sigma)(\mathcal{D}, \theta, Q) \\ &= \phi_1(\mathcal{D}, \theta_1) \left[\sum_z \phi_2(\mathcal{D}, \theta_1, z) \right] \phi_3(\mathcal{D}, \theta_1, \theta_3) p(z \mid \mathcal{D}, \theta_1) p(Q \mid \mathcal{D}, \theta) \quad \text{from Figure 13,} \\ &= \left[\sum_z \phi_1(\mathcal{D}, \theta_1) \phi_2(\mathcal{D}, \theta_1, z) \phi_3(\mathcal{D}, \theta_1, \theta_3) \right] p(z \mid \mathcal{D}, \theta_1) p(Q \mid \mathcal{D}, \theta) \quad \text{by the distributive} \\ &\quad \text{law,} \\ &\propto \left[\sum_z p(\mathcal{D}, \theta_1, z, \theta_2) \right] p(z \mid \mathcal{D}, \theta_1) p(Q \mid \mathcal{D}, \theta) \quad \text{by Theorem 1} \\ &= p(\mathcal{D}, \theta_1, \theta_2) p(z \mid \mathcal{D}, \theta_1) p(Q \mid \mathcal{D}, \theta) \quad \text{and Lemma 10,} \\ &= p(\mathcal{D}, \theta_1, \theta_2) p(z \mid \mathcal{D}, \theta_1, \theta_3) p(Q \mid \mathcal{D}, \theta) \quad \text{marginalisation of } z, \\ &= p(\mathcal{D}, \theta_1, \theta_2) p(z \mid \mathcal{D}, \theta_1, \theta_3) p(Q \mid \mathcal{D}, \theta) \quad \text{by } z \perp \theta_3 \mid \theta_1 \\ &\quad \text{(Theorem 3),} \\ &= p(\mathcal{D}, \theta, Q) \quad \text{by the chain rule} \\ &\propto \llbracket S \rrbracket_p(\sigma)(\mathcal{D}, \theta, Q) \quad \text{for probability,} \end{aligned}$$

Together, this gives us $\llbracket S \rrbracket = \llbracket S' \rrbracket$ (up to S' creating one new fresh variable f). \square

Fig. 13. Diagrammatic proof of semantic preservation of \xrightarrow{z} .

B EXAMPLES

B.1 Sprinkler

Often, beginners are introduced to probabilistic modelling through simple, discrete variable examples, as they are more intuitive to reason about and often have analytical solutions. Unfortunately, one cannot express such examples directly in PPLs that do not support discrete parameters. One well-known discrete variable example, often used in tutorials on probabilistic modelling, is the “Sprinkler” example. It models the relationship between cloudy weather, whether it rains, whether the garden sprinkler is on, and the wetness of the grass. In Figure 14, we show a version of the sprinkler model written in SlicStan with discrete parameters (left) and the marginalisation part of its corresponding transformed version (right).

As cloudy \perp wet | sprinkler, rain, we do not need to include wet in the elimination of cloudy, and the new factor is computed for different values of only sprinkler and rain (lines 2–6). The rest of the variables are eliminated one-by-one, involving all remaining variables (lines 7–15).

The snippet of the SlicStan code generated by our transformation is an exact implementation of the variable elimination algorithm for this model. This not only facilitates a platform for learning probabilistic programming using standard introductory models, but it can also be a useful tool for learning concepts such as marginalisation, conditional independence, and exact inference methods.

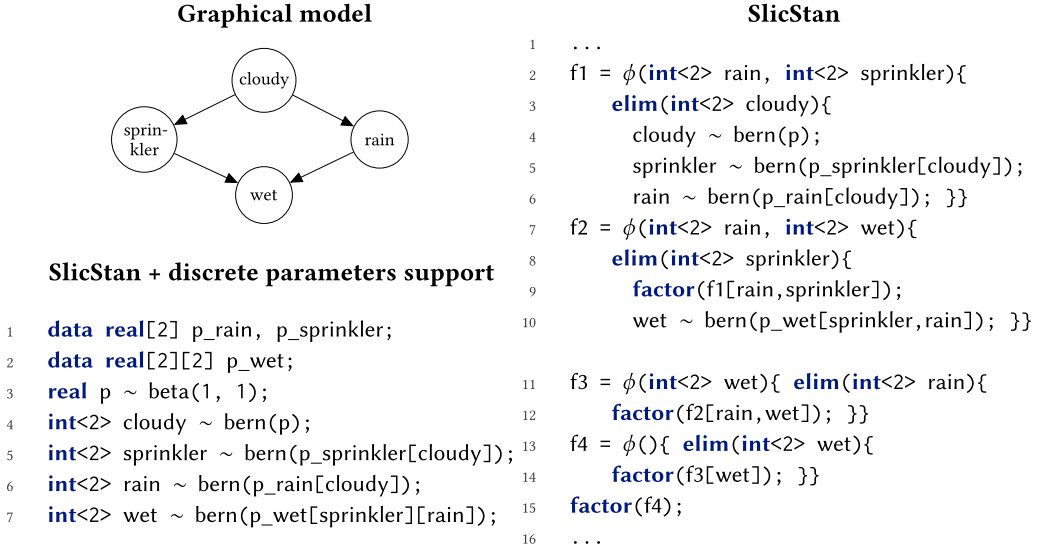


Fig. 14. The “Sprinkler” example.

B.2 Soft-K-means Model

In Figure 15, we present the standard soft-k-means clustering model as it is written in SlicStan with support for discrete model parameters (left). The right column shows the resulting code that our program transformation generates. This code consists of plain SlicStan code and no support for discrete model parameters is needed to perform inference on it.

The model can be used for (softly) dividing N data points y in D -dimensional Euclidean space into K clusters that have means μ and probability π .

SlicStan + discrete	SlicStan
<pre> data int D; data int K; data real[K] pi; data real N = 3; data real[D][N] y; real[D][K] mu; for(d in 1 : D) { for(k in 1 : K){ mu[d][k] ~ normal(0, 1); } int<K> z1 ~ categorical(pi); int<K> z2 ~ categorical(pi); int<K> z3 ~ categorical(pi); for(d in 1 : D) { y[d][1] ~ normal(mu[d][z1], 1); y[d][2] ~ normal(mu[d][z2], 1); y[d][3] ~ normal(mu[d][z3], 1); } </pre>	<pre> ... for(d in 1:D){ for(k in 1:K){ mu[d,k] ~ normal(0, 1);}} factor(elim(int<K> z1){ z1 ~ categorical(pi); for(data int d in 1:D){ y[d,1] ~ normal(mu[d,z1], 1);}}); factor(elim(int<K> z2){ z2 ~ categorical(pi); for(data int d in 1:D){ y[d,2] ~ normal(mu[d,z2], 1);}}); factor(elim(int<K> z3){ z3 ~ categorical(pi); for(data int d in 1:D){ y[d,3] ~ normal(mu[d,z3], 1);}}); gen(int z3){ z3 ~ categorical(pi); for(data int d in 1:D){ y[d,3] ~ normal(mu[d,z3], 1);}} gen(int z2){ z2 ~ categorical(pi); for(data int d in 1:D){ y[d,2] ~ normal(mu[d,z2], 1);}} gen(int z1){ z1 ~ categorical(pi); for(data int d in 1:D){ y[d,1] ~ normal(mu[d,z1], 1);}} </pre>

Fig. 15. Soft K -means.

B.3 A Causal Inference Example

The question of how to adapt PPLs to causal queries has been recently gaining popularity. One way to express interventions and reason about causality, is to assume a discrete variable specifying the direction (or absence of) causal relationship and specify different behaviour for each case using if statements [Winn 2012]. We show a simple causal inference example (Figure 16) written in SlicStan with direct support for discrete parameters (left) and the code that our transformation generates (right) on which we can perform inference using a combination of, e.g., HMC and ancestral sampling.

This model can be read as follows: Assume that we are in a situation where we want to answer a causal question. We want to answer this question based on N paired observations of A and B , in some of which we might have intervened (doB). Our model proceeds by drawing a (prior) probability that A causes B from a beta distribution, and then specifying A and B for different scenarios (intervention, A causes B and no intervention, B causes A and no intervention) using conditional statements.

SlicStan + discrete	SlicStan
<pre> data real q; data int N; data int[N] A, B, doB; data real prob_intervention; real pAcausesB ~ beta(1, 1); int<2> AcausesB ~ bernoulli(pAcausesB); for (n in 1:N) if(doB[n] > 0) B[n] ~ bernoulli(prob_intervention); if (AcausesB > 1){ for (n in 1:N){ A[n] ~ bernoulli(0.5); if (doB[n] < 1){ if (A[n] > 0) { B[n] ~ bernoulli(q); } else { B[n] ~ bernoulli(1 - q); } } } } else { for (n in 1:N){ if (doB[n] < 1){ B[n] ~ bernoulli(0.5); } if (B[n] > 0){ A[n] ~ bernoulli(q); } else { A[n] ~ bernoulli(1 - q); } } } </pre>	<pre> data real q; data int N; data int[N] A, B, doB; data real prob_intervention; real pAcausesB ~ beta(1, 1); for(data int n in 1:N) if(doB[n] > 0) B[n] ~ bernoulli(prob_intervention); factor(elim(int<2> AcausesB){ AcausesB ~ bernoulli(pAcausesB); if(AcausesB > 1){ for(data int n in 1:N){ A[n] ~ bernoulli(0.5); if(doB[n] < 1){ if(A[n] > 0){B[n] ~ bernoulli(q);} else{ B[n] ~ bernoulli(1 - q); } } } } } else{ for(data int n in 1:N){ if(doB[n] < 1){ B[n] ~ bernoulli(0.5); } if(B[n] > 0){ A[n] ~ bernoulli(q); } else{ A[n] ~ bernoulli(1 - q); } } }}}); </pre>

Fig. 16. A causal inference example.

ACKNOWLEDGMENTS

We thank Vikash Mansinghka for suggesting the outlier detection example, which we used for evaluation, as well as Lawrence Murry for clarifying the behaviour of Birch, and anonymous reviewers whose helpful suggestions improved the article.

REFERENCES

- Martín Abadi, Anindya Banerjee, Nevin Heintze, and Jon G. Riecke. 1999. A core calculus of dependency. In *Symposium on Principles of Programming Languages*. ACM, 147–160.
- Eyal Amir. 2010. Approximation algorithms for treewidth. *Algorithmica* 56, 4 (2010), 448–479.
- Torben Amtoft and Anindya Banerjee. 2020. A theory of slicing for imperative probabilistic programs. *ACM Trans. Program. Lang. Syst.* 42, 2 (2020), 1–71.
- Stefan Arnborg, Derek G. Corneil, and Andrzej Proskurowski. 1987. Complexity of finding embeddings in a k-tree. *SIAM J. Algeb. Discr. Meth.* 8, 2 (1987), 277–284.
- Jialu Bao, Simon Docherty, Justin Hsu, and Alexandra Silva. 2021. A bunched logic for conditional independence. In *36th Annual ACM/IEEE Symposium on Logic in Computer Science*. IEEE, 1–14.

- Gilles Barthe, Justin Hsu, and Kevin Liao. 2019. A probabilistic separation logic. *Proc. ACM on Program. Lang.* 4, POPL (2019), 1–30.
- Michael Betancourt and Mark Girolami. 2015. Hamiltonian Monte Carlo for hierarchical models. *Curr. Trends Bayesian Methodol. Applic.* 79 (2015), 30.
- Craig Boutilier, Nir Friedman, Moises Goldszmidt, and Daphne Koller. 1996. Context-specific independence in Bayesian networks. In *12th International Conference on Uncertainty in Artificial Intelligence*. 115–123.
- Bob Carpenter, Andrew Gelman, Matthew Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. 2017. Stan: A probabilistic programming language. *J. Statist. Softw.* 76, 1 (2017), 1–32. DOI : <https://doi.org/10.18637/jss.v076.i01>
- Marco F. Cusumano-Towner, Feras A. Saad, Alexander K. Lew, and Vikash K. Mansinghka. 2019. Gen: A general-purpose probabilistic programming system with programmable inference. In *40th ACM SIGPLAN Conference on Programming Language Design and Implementation*. Association for Computing Machinery, New York, NY, 221–236. DOI : <https://doi.org/10.1145/3314221.3314642>
- Luis Damiano, Brian Peterson, and Michael Weylandt. 2018. A tutorial on hidden Markov models using Stan. *StanCon* (2018). DOI : <https://doi.org/10.5281/zenodo.1284341>.
- Leonardo De Moura and Nikolaj Bjørner. 2008. Z3: An efficient SMT solver. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 337–340.
- Paul Feautrier. 1992. Some efficient solutions to the affine scheduling problem. Part II. Multidimensional time. *Int. J. Parallel Program.* 21, 6 (1992), 389–420.
- Brendan J. Frey. 2002. Extending factor graphs so as to unify directed and undirected graphical models. In *19th Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann Publishers Inc., San Francisco, CA, 257–264.
- Hong Ge, Kai Xu, and Zoubin Ghahramani. 2018. Turing: A language for flexible probabilistic inference. In *International Conference on Artificial Intelligence and Statistics*. 1682–1690. DOI : <http://proceedings.mlr.press/v84/ge18b.html>.
- Timon Gehr, Sasa Misailovic, and Martin Vechev. 2016. PSI: Exact symbolic inference for probabilistic programs. In *International Conference on Computer-aided Verification*. Springer, 62–83.
- Timon Gehr, Samuel Steffen, and Martin Vechev. 2020. λPSI: Exact inference for higher-order probabilistic programs. In *41st ACM SIGPLAN Conference on Programming Language Design and Implementation*. 883–897.
- Andrew Gelman, Daniel Lee, and Jiqiang Guo. 2015. Stan: A probabilistic programming language for Bayesian inference and optimization. *J. Educ. Behav. Statist.* 40, 5 (2015), 530–543.
- Andrew D. Gordon, Claudio V. Russo, Marcin Szymczak, Johannes Borgström, Nicolas Rolland, Thore Graepel, and Daniel Tarlow. 2015. Probabilistic programs as spreadsheet queries. In *European Symposium on Programming Languages and Systems (Lecture Notes in Computer Science, Vol. 9032)*. Springer, 1–25.
- Maria I. Gorinova, Andrew D. Gordon, and Charles Sutton. 2019. Probabilistic programming with densities in SlicStan: Efficient, flexible, and deterministic. *Proce. ACM Program. Lang.* 3, POPL (2019), 35.
- Andreas Griewank and Andrea Walther. 2008. *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation*. SIAM.
- Matthew D. Hoffman and Andrew Gelman. 2014. The No-U-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *J. Mach. Learn. Res.* 15, 1 (2014), 1593–1623.
- Steven Holtzen, Guy Van den Broeck, and Todd Millstein. 2020. Dice: Compiling discrete probabilistic programs for scalable inference. *arXiv preprint arXiv:2005.09089* (2020).
- Uffe Kjærulff. 1990. Triangulation of graphs-Algorithms giving small total state space. Research Report R 90-09. Department of Mathematics and Computer Science, Aalborg University, Denmark.
- Daphne Koller and Nir Friedman. 2009. *Probabilistic Graphical Models: Principles and Techniques*. The MIT Press.
- Elisabet Lobo-Vesga, Alejandro Russo, and Marco Gaboardi. 2020. A programming framework for differential privacy with accuracy concentration bounds. In *IEEE Symposium on Security and Privacy*. IEEE, 411–428.
- Vikash K. Mansinghka, Ulrich Schaechtle, Shivam Handa, Alexey Radul, Yutian Chen, and Martin Rinard. 2018. Probabilistic programming with programmable inference. In *39th ACM SIGPLAN Conference on Programming Language Design and Implementation*. ACM, New York, NY, 603–616. DOI : <https://doi.org/10.1145/3192366.3192409>
- Tom Minka and John Winn. 2009. Gates. In *Advances in Neural Information Processing Systems*, D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou (Eds.), Vol. 21. Curran Associates, Inc. Retrieved from <https://proceedings.neurips.cc/paper/2008/file/4b0a59ddf11c58e7446c9df0da541a84-Paper.pdf>.
- T. Minka, J. M. Winn, J. P. Guiver, S. Webster, Y. Zaykov, B. Yangel, A. Spengler, and J. Bronskill. 2014. Infer.NET 2.6. Microsoft Research Cambridge. Retrieved from <http://research.microsoft.com/infernet>.
- Dave Moore and Maria I. Gorinova. 2018. Effect handling for composable program transformations in Edward2. In *International Conference on Probabilistic Programming*. <https://arxiv.org/abs/1811.06150>.
- Kevin P. Murphy. 2012. *Machine Learning: A Probabilistic Perspective*. The MIT Press.

- Lawrence M. Murray, Daniel Lundén, Jan Kudlicka, David Broman, and Thomas B. Schön. 2018. Delayed sampling and automatic Rao-Blackwellization of probabilistic programs. In *21st International Conference on Artificial Intelligence and Statistics*. 1037–1046. Retrieved from <http://proceedings.mlr.press/v84/murray18a.html>.
- Lawrence M. Murray and Thomas B. Schön. 2018. Automated learning with a probabilistic programming language: Birch. *Ann. Rev. Contr.* 46 (2018), 29–43.
- Praveen Narayanan, Jacques Carette, Wren Romano, Chung-chieh Shan, and Robert Zinkov. 2016. Probabilistic inference by program transformation in Hakaru (system description). In *Functional and Logic Programming*, Oleg Kiselyov and Andy King (Eds.). Springer International Publishing, Cham, 62–79.
- Radford M. Neal et al. 2011. MCMC using Hamiltonian dynamics. *Handb. Mark. Chain Monte Carlo* 2, 11 (2011).
- Akihiko Nishimura, David Dunson, and Jianfeng Lu. 2017. Discontinuous Hamiltonian Monte Carlo for sampling discrete parameters. *arXiv preprint arXiv:1705.08510* (2017).
- Fritz Obermeyer, Eli Bingham, Martin Jankowiak, Justin Chiu, Neeraj Pradhan, Alexander Rush, and Noah Goodman. 2019. Tensor variable elimination for plated factor graphs. In *International Conference on Machine Learning*.
- A. Pakman and L. Paninski. 2013. Auxiliary-variable exact Hamiltonian Monte Carlo samplers for binary distributions. In *International Conference on Advances in Neural Information Processing Systems*.
- Du Phan, Neeraj Pradhan, and Martin Jankowiak. 2019. Composable effects for flexible and accelerated probabilistic programming in NumPyro. *arXiv preprint arXiv:1912.11554* (2019).
- Benjamin C. Pierce and David N. Turner. 2000. Local type inference. *ACM Trans. Program. Lang. Syst.* 22, 1 (2000), 1–44.
- Lawrence R. Rabiner. 1989. A tutorial on hidden Markov models and selected applications in speech recognition. *Proc. IEEE* 77, 2 (1989), 257–286.
- John Salvatier, Thomas V. Wiecki, and Christopher Fonnesbeck. 2016. Probabilistic programming in Python using PyMC3. *PeerJ Comput. Sci.* 2 (2016), e55.
- Stan Development Team. 2019a. Stan language reference manual. Version 2.26.0. Retrieved from <http://mc-stan.org>.
- Stan Development Team. 2019b. Stan user’s guide. Version 2.26.0. Retrieved from <http://mc-stan.org>.
- Dustin Tran, Matthew D. Hoffman, Srinivas Vasudevan, Christopher Suter, Dave Moore, Alexey Radul, Matthew Johnson, and Rif A. Saurous. 2018. Edward2: Simple, distributed, accelerated. Retrieved from https://github.com/tensorflow/probability/tree/master/tensorflow_probability/python/edward2.
- Uber AI Labs. 2017. Pyro: A deep probabilistic programming language. Retrieved from <http://pyro.ai/>.
- Dennis M. Volpano, Cynthia E. Irvine, and Geoffrey Smith. 1996. A sound type system for secure flow analysis. *J. Comput. Secur.* 4, 2/3 (1996), 167–188.
- John Winn. 2012. Causality with gates. In *International Conference on Artificial Intelligence and Statistics*. 1314–1322. Retrieved from <http://proceedings.mlr.press/v22/winn12.html>.
- Frank Wood, Jan Willem van de Meent, and Vikash Mansinghka. 2014. A new approach to probabilistic programming inference. In *17th International Conference on Artificial Intelligence and Statistics*. 1024–1032.
- Nevin Lianwen Zhang and David Poole. 1994. A simple approach to Bayesian network computations. In *Conference of the Canadian Society for Computational Studies of Intelligence*. Canadian Information Processing Society, 171–178.
- Yichuan Zhang, Zoubin Ghahramani, Amos J. Storkey, and Charles A. Sutton. 2012. Continuous relaxations for discrete Hamiltonian Monte Carlo. In *Advances in Neural Information Processing Systems*, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger (Eds.). Curran Associates, Inc., 3194–3202. Retrieved from <http://papers.nips.cc/paper/4652-continuous-relaxations-for-discrete-hamiltonian-monte-carlo.pdf>.
- Guangyao Zhou. 2020. Mixed Hamiltonian Monte Carlo for mixed discrete and continuous variables. In *International Conference on Advances in Neural Information Processing Systems*.
- Yuan Zhou, Bradley J. Gram-Hansen, Tobias Kohn, Tom Rainforth, Hongseok Yang, and Frank Wood. 2019. LF-PPL: A low-level first order probabilistic programming language for non-differentiable models. In *International Conference on Artificial Intelligence and Statistics*. Retrieved from <http://proceedings.mlr.press/v89/zhou19b.html>.

Received April 2021; revised August 2021; accepted September 2021