



Probabilistic programs are getting big, and their verification at scale requires modularity ①.

Parallels between probability and mutable state suggest the use of separation logic ②.

① Statistical independence is a fundamental modularity principle: probabilistic reasoning frequently proceeds by decomposition into independent pieces.

② Both sampling and allocation are generative: sampling yields a random variable independent of previous ones, just as allocation yields a fresh chunk of memory.

Some separation logics support independence substructurally ③, but they overapproximate ④.

③ In these logics, propositions are sets of distributions on stores, and $\mu \vDash P * Q$ if μ factors into independent distributions μ_1 and μ_2 on stores with disjoint domain.

Some common independences are not expressible:

④ $X = \text{flip } 1/2; Y = \text{flip } 1/2; Z = X + Y$
 $\{(Z - X) * (Z - Y)\}$

This postcondition is unprovable because Z appears twice.

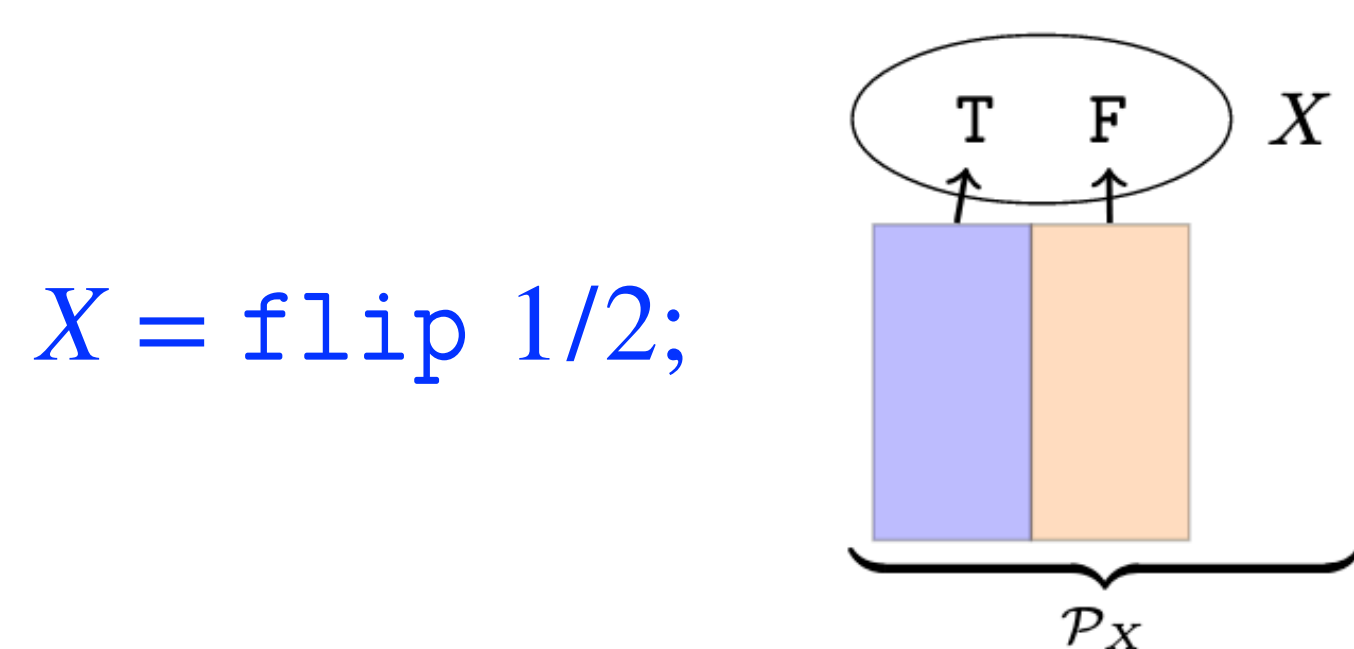
We present a new model of separation logic where separating conjunction is interpreted by the *independent combination* of probability spaces ⑤. This novel combining operation is the probabilistic analogue of disjoint union of heaps ⑥.

⑤ Fix a sample space Ω . The *independent combination* of two probability spaces (\mathcal{F}, μ) and (\mathcal{G}, ν) is $(\mathcal{F}, \mu) \cdot (\mathcal{G}, \nu) = (\sigma(\mathcal{F}, \mathcal{G}), \rho)$ where $\rho(F \cap G) = \mu(F)\nu(G)$ for all $F \in \mathcal{F}, G \in \mathcal{G}$.

Independent combination forms a Kripke resource monoid, giving the expected interpretation of separating conjunction:


⑥ $(\mathcal{E}, \rho) \vDash P * Q \iff$ there exists (\mathcal{F}, μ) and (\mathcal{G}, ν) with $(\mathcal{F}, \mu) \vDash P$ and $(\mathcal{G}, \nu) \vDash Q$ and $(\mathcal{E}, \rho) = (\mathcal{F}, \mu) \cdot (\mathcal{G}, \nu)$.

Using independent combination, one can read probabilistic programs operationally: sampling literally allocates probability spaces, and conditioning performs destructive update ⑧.




Shaded regions denote events, and their areas denote probabilities.

The 1st **flip** allocates a probability space \mathcal{P}_X with blue-orange σ -algebra and a \mathcal{P}_X -measurable random variable X with distribution $\text{Ber } 1/2$.



The 2nd **flip** allocates \mathcal{P}_Y with the dotted-dashed σ -algebra and a \mathcal{P}_Y -measurable variable Y , producing the independent combination $\mathcal{P}_X \cdot \mathcal{P}_Y$.



Finally, **observe** destructively updates the measure stored in \mathcal{P}_X so that $P(X = T) = 1$.

This new separation logic precisely captures independence ⑦, enjoys a completely standard frame rule, and easily supports continuous random variables. Using disintegration theory, we extend the base logic with a *conditioning modality* ⑧, and use this modal logic to verify a challenging randomized algorithm ⑨.

⑦ $(\mathcal{F}, \mu) \vDash X_1 * \dots * X_n$ iff X_1, \dots, X_n mutually independent.

$D_{x \leftarrow X} P$ says P holds conditional on $X = x$ for all x . Propositions have intuitive "conditional" readings

⑧ under D . For example, $D_{x \leftarrow X}(Y * Z)$ expresses conditional independence of Y and Z given X .

⑨ Sampling from a finite distribution can be done in constant space given access to continuous variables. The correctness proof uses D to condition on a continuous variable and a derived rule expressing the law of total expectation. A crucial step exploits independences automatically preserved through each loop iteration by the frame rule.