

Learning Task-General Representations with Generative Neuro-Symbolic Modeling

Reuben Feinman & Brenden Lake


ICLR 2021

Structure and statistics in conceptual representations

Ingredients (parts)



Correlations



substitutions



co-occurences

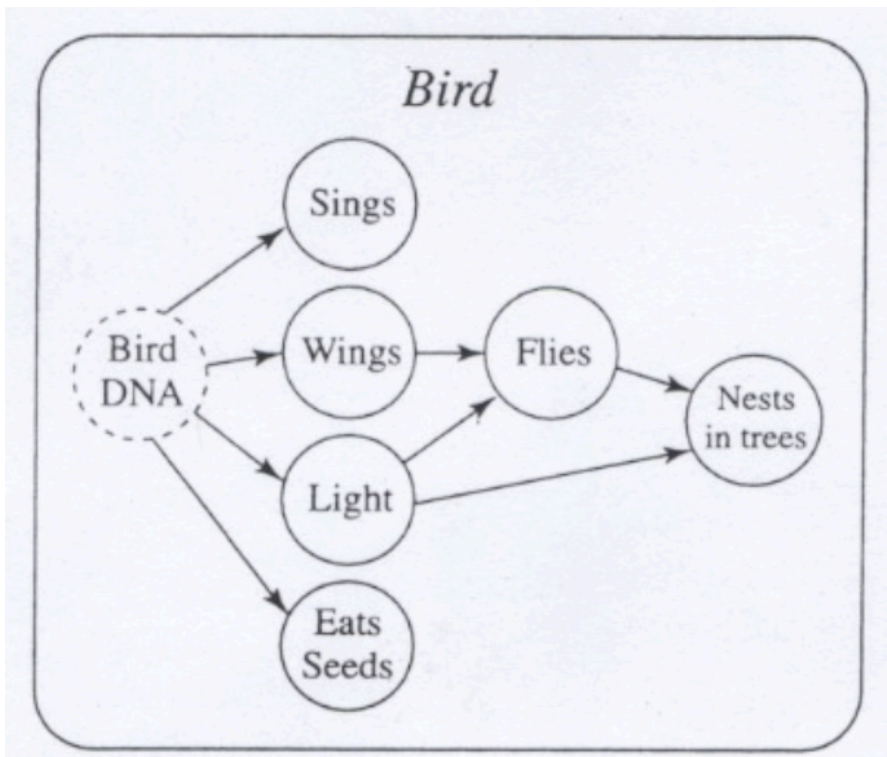
Causal processes



Modeling structure and statistics

Structured knowledge (symbolic models)


Causal-Model Theory
(Rehder, 2007)



Bayesian Program Learning
(Lake et al., 2015)

```


procedure GENERATE TYPE
   $\kappa \leftarrow P(\kappa)$  ▷ Sample number of parts
  for  $i = 1 \dots \kappa$  do
     $n_i \leftarrow P(n_i | \kappa)$  ▷ Sample number of sub-parts
    for  $j = 1 \dots n_i$  do
       $s_{ij} \leftarrow P(s_{ij} | s_{i(j-1)})$  ▷ Sample sub-part sequence
    end for
     $R_i \leftarrow P(R_i | S_1, \dots, S_{i-1})$  ▷ Sample relation
  end for
   $\psi \leftarrow \{\kappa, R, S\}$ 
  return @GENERATE TOKEN( $\psi$ ) ▷ Return program
  
```




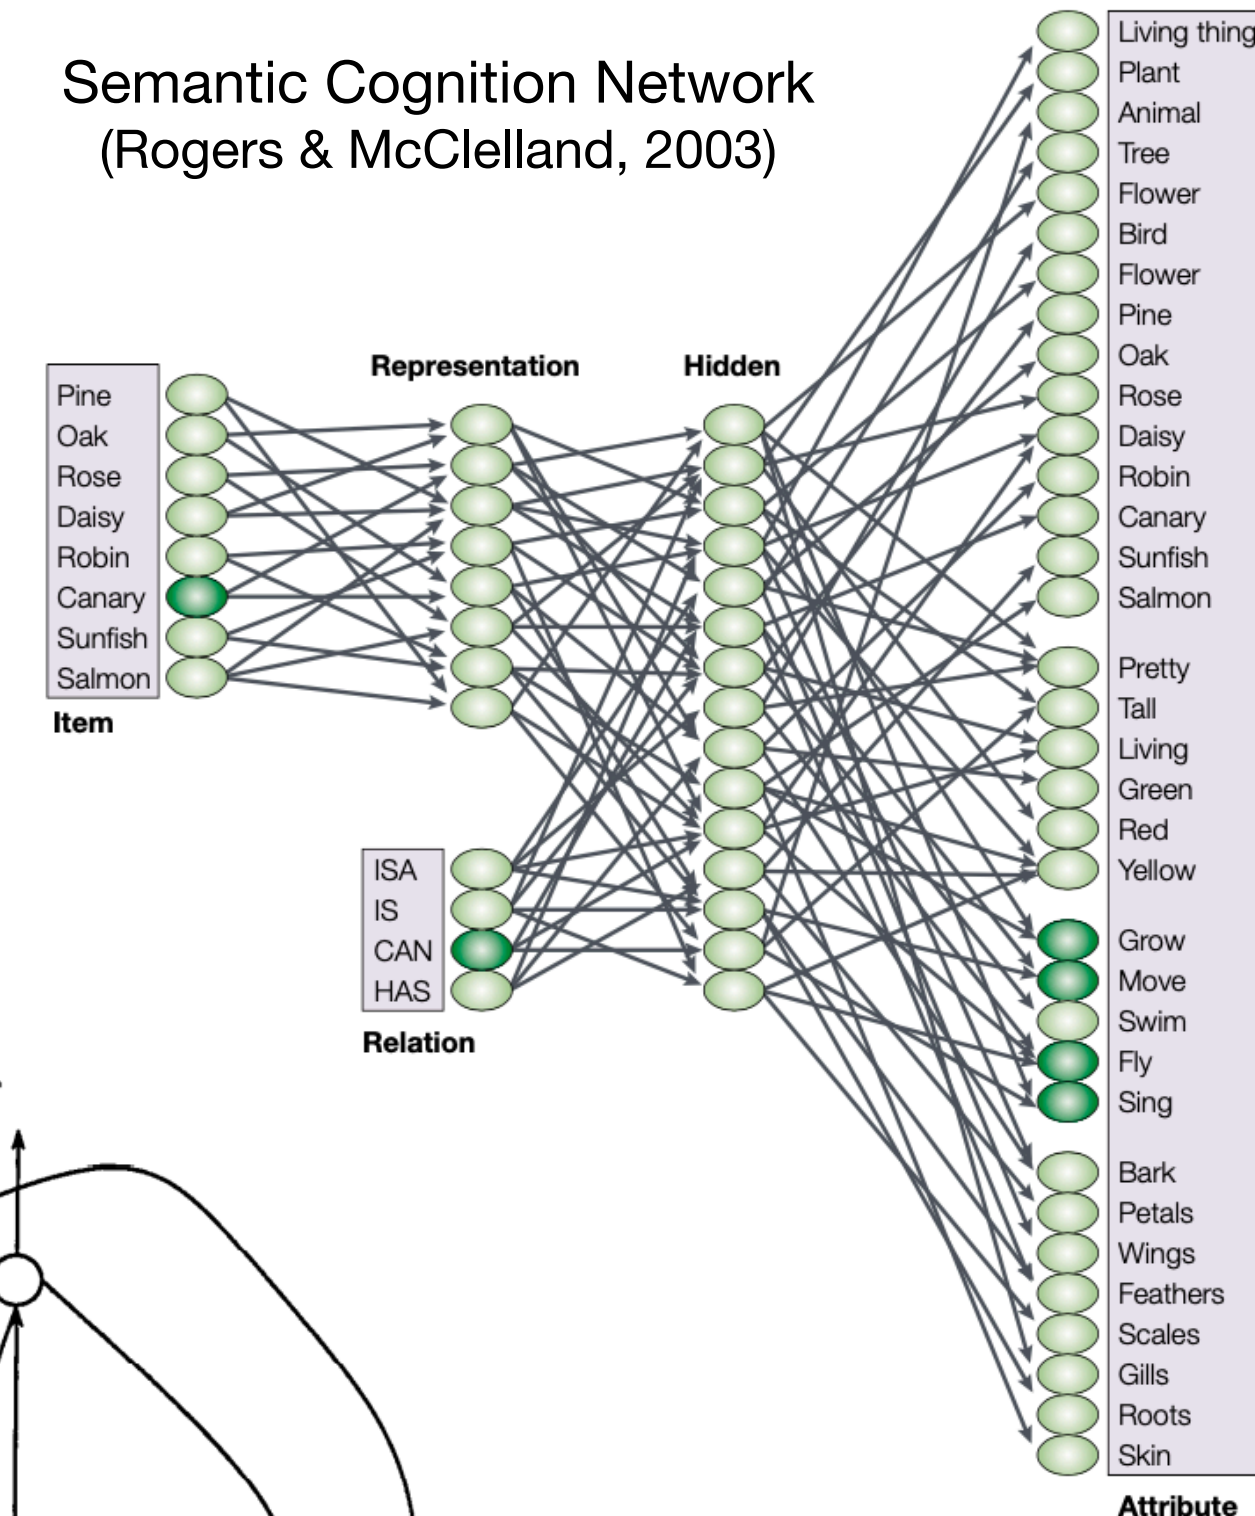
Structural Forms
(Kemp & Tenenbaum, 2008)

Emergent "statistical" knowledge (neural networks)

ALCOVE
(Kruschke, 1992)



Semantic Cognition Network
(Rogers & McClelland, 2003)



Finding Structure in Time
(Elman, 1990)

Proposal: Generative Neuro-Symbolic (GNS) Modeling

- Goal: model the *compositional* and *causal* structure in how concepts are formed, while simultaneously modeling nonparametric statistical relationships
- Proposal: probabilistic programs with neural network sub-routines

procedure GENERATECONCEPT

$M \leftarrow 0$

▷ Initialize memory state

while *True* **do**

$x_i, r_i \sim p(x, r \mid M)$

▷ Sample part and relation from neural net

$M \leftarrow f_{render}(x_i, r_i, M)$

▷ Render part to memory (differentiable)

$v_i \sim p(v \mid M)$

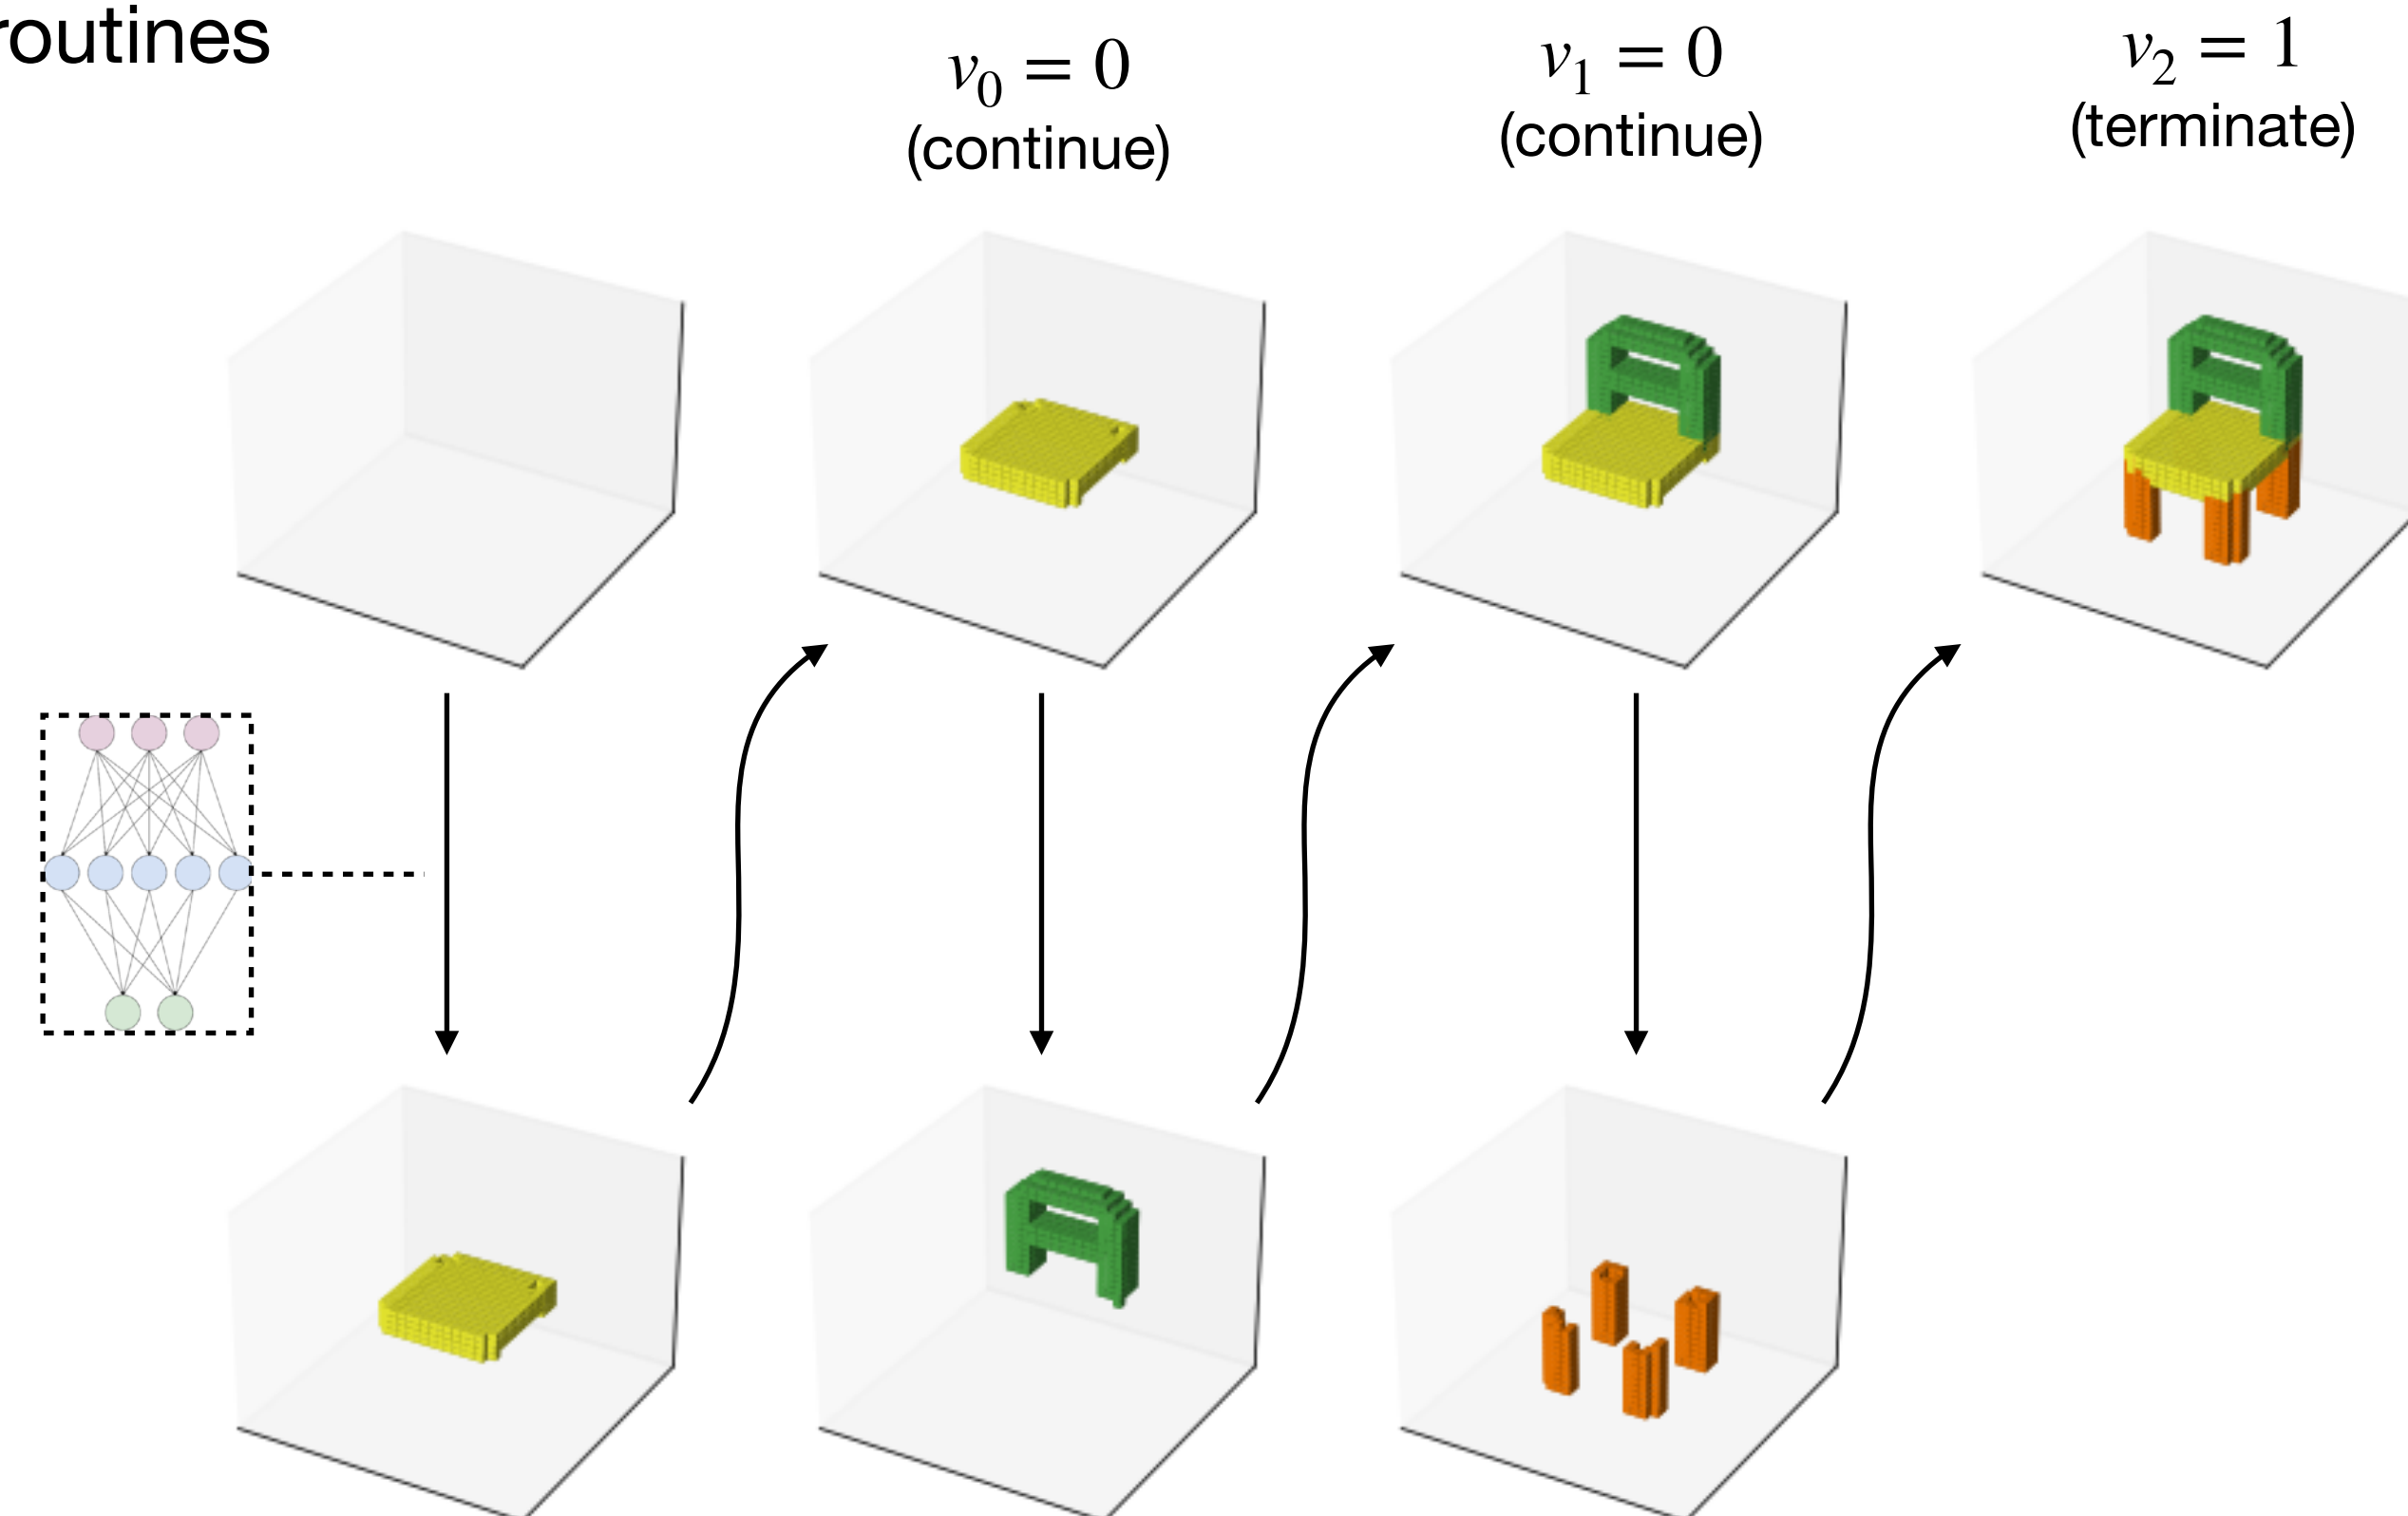
▷ Sample termination indicator

if v_i **then**

break

return $\{X, R\}$


▷ Return concept type



Omniglot concept learning challenge

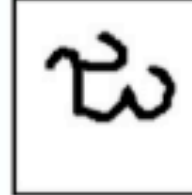
(Lake et al., 2015)

(a) One-shot classification



Where is another?

ग	प	म	र	च
क	ख	श	ष	झ
ल	व	स	ह	य
न	य	म	न	ः



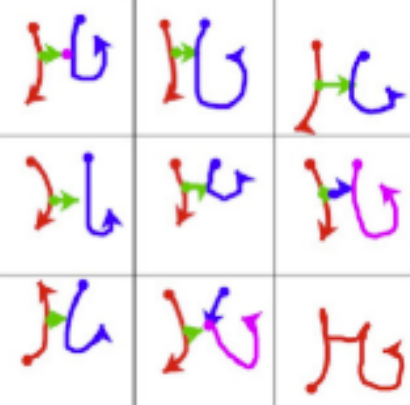
Where is another?

ल	व	ल	व	ड
क	ख	ग	ड	ध
श	च	छ	ट	ठ
र	य	ल	क	ड

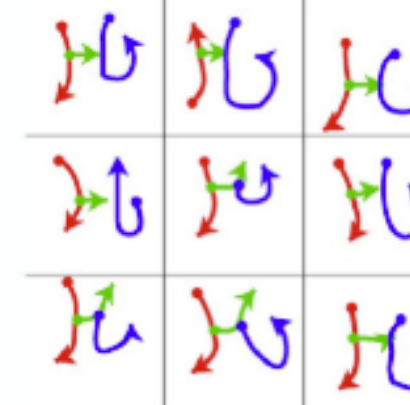
(b) Parsing

श	ड	ड
श	ड	ड
श	ड	ड

Human parses




Machine parses



stroke order — 1 — 2 — 3 — 4 — 5

(c) Generating new exemplars



Human or Machine?

1

श	ड	ड
श	ड	ड
श	ड	ड

2

श	ड	ड
श	ड	ड
श	ड	ड

(d) Generating new concepts (from type)

Alphabet

ग	ख	ग	ख	ड
र	ड	ग	ड	ः

Human or Machine?

1

ग	ख
र	ड

2

ग	ख
र	ड

(e) Generating new concepts (unconstrained)

Human or Machine?

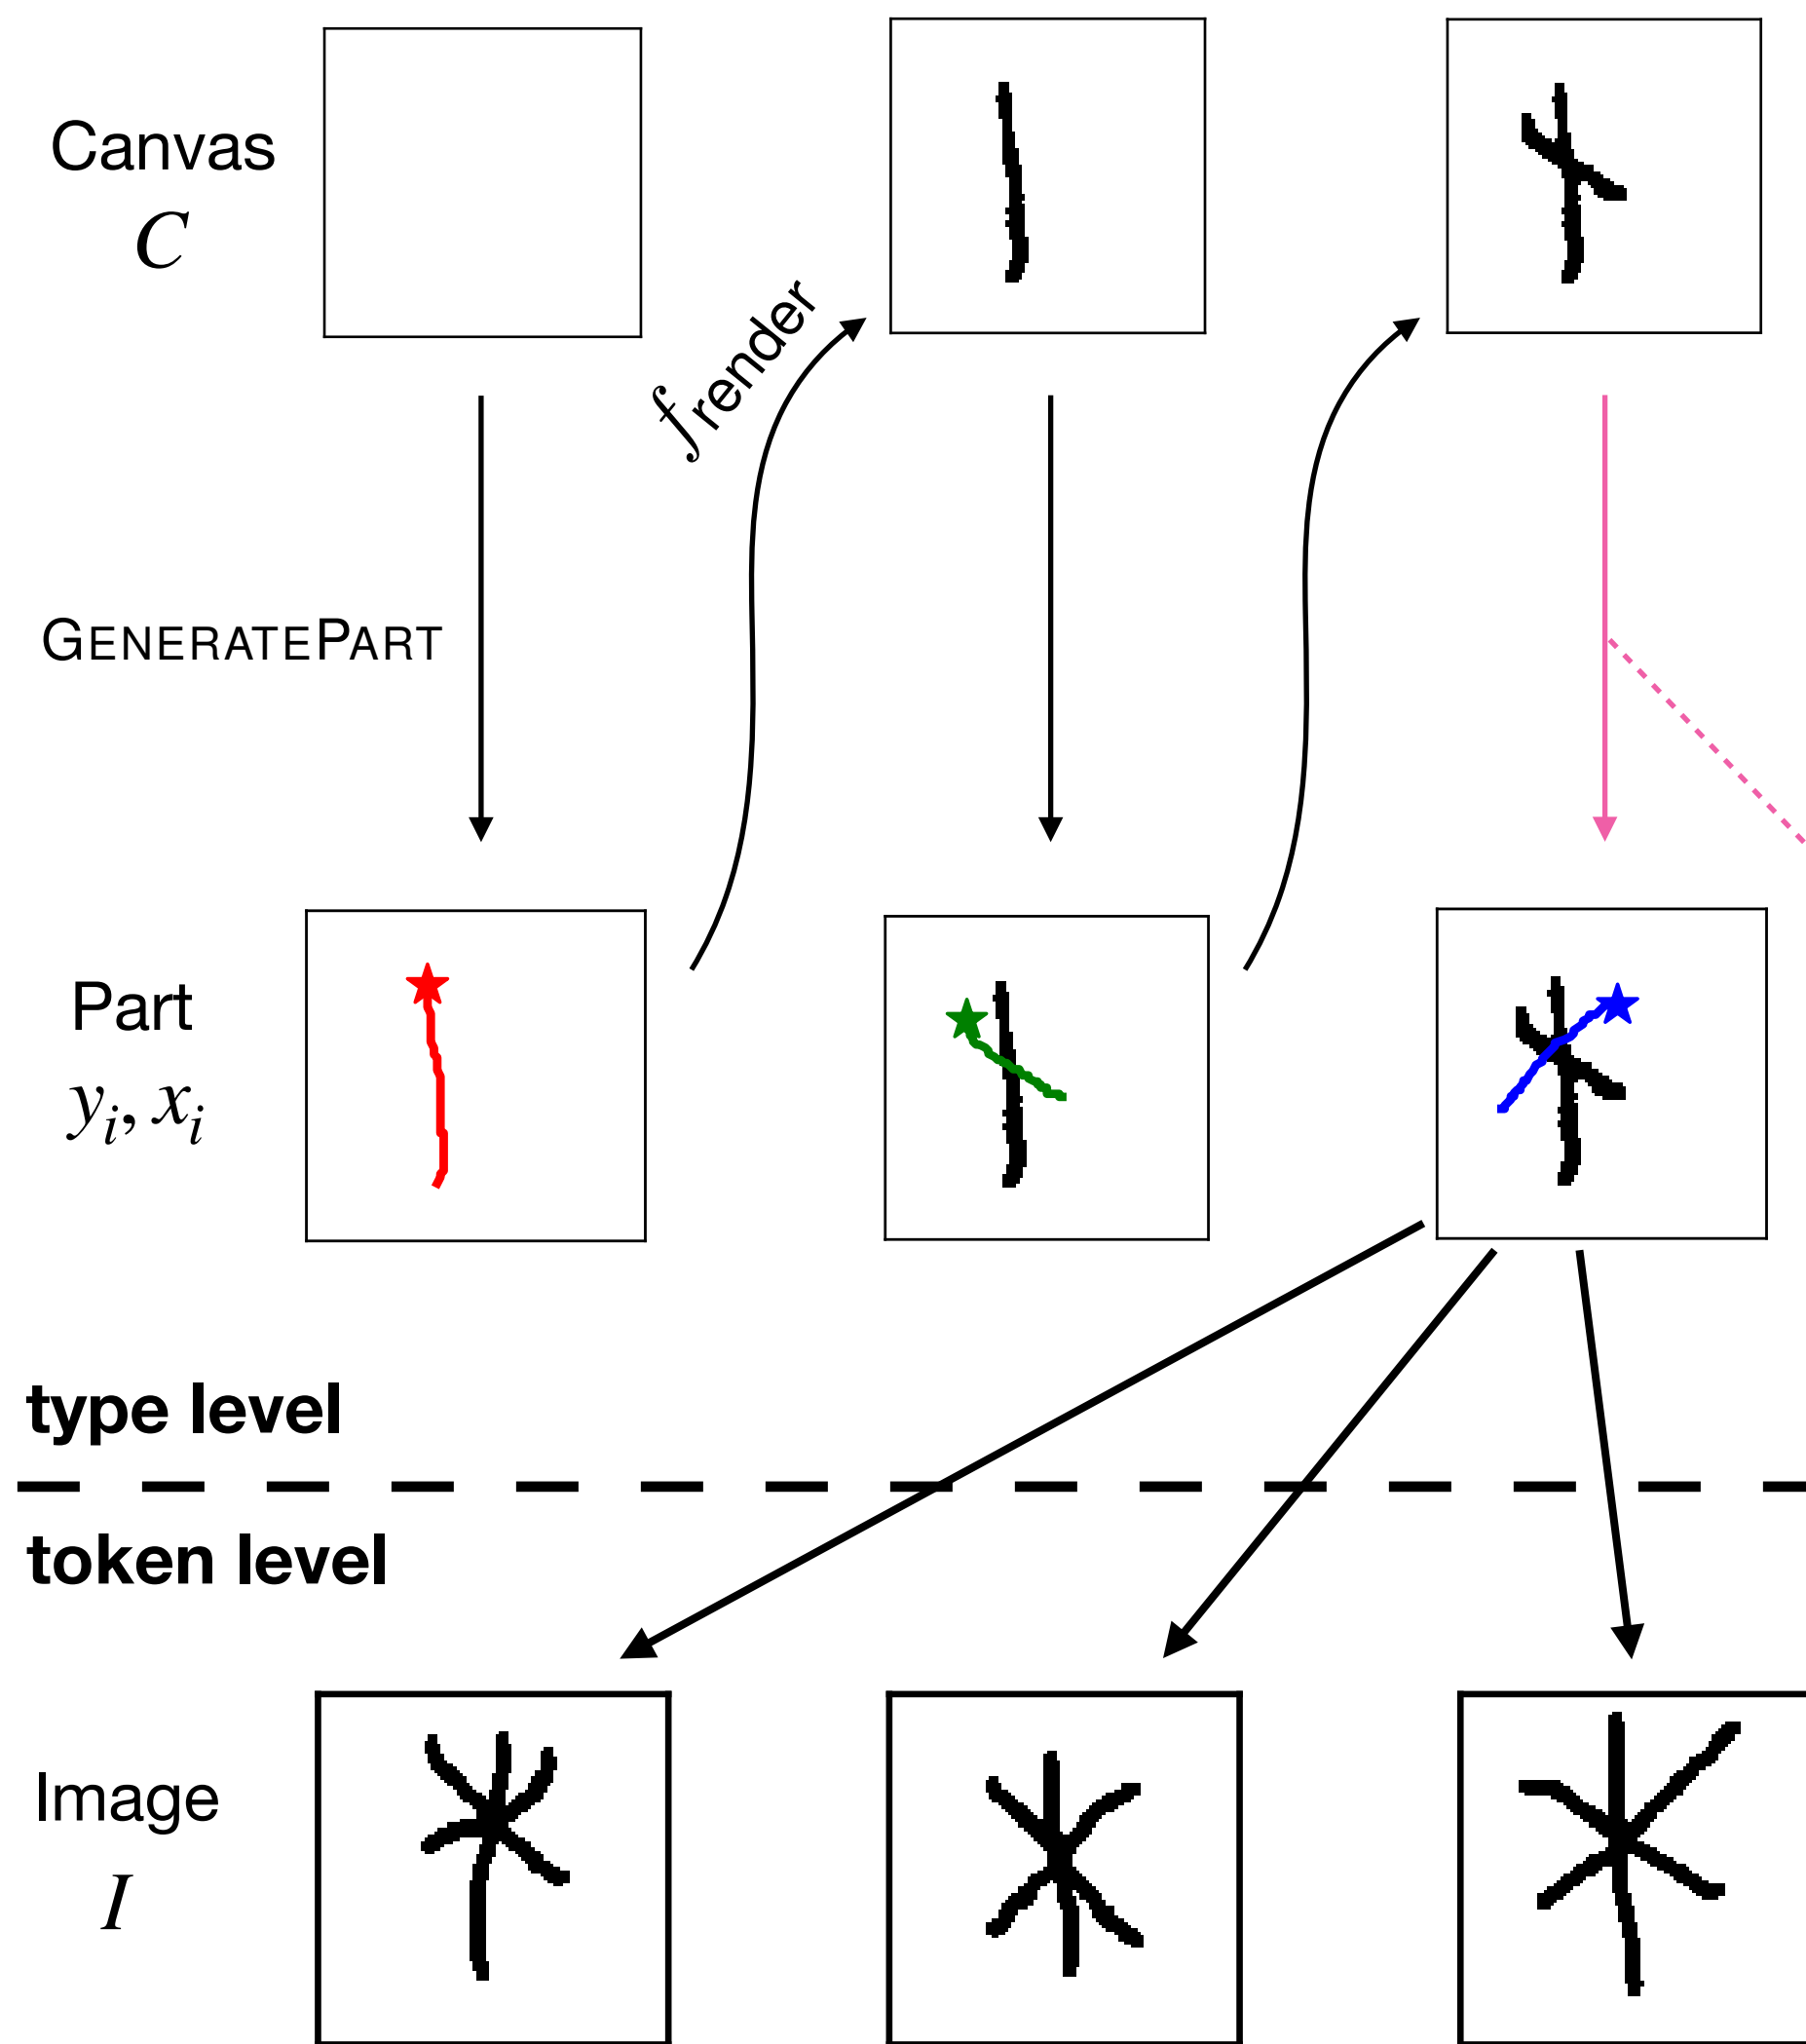
1

श	ड
श	ड

2

श	ड
श	ड

GNS model of character concepts



procedure GENERATETYPE

$C \leftarrow 0$

while *true* **do**

$[y_i, x_i] \leftarrow \text{GENERATEPART}(C)$

$C \leftarrow f_{\text{render}}(y_i, x_i, C)$

$v_i \sim p(v \mid C)$

if v_i **then**

break

$\psi \leftarrow \{\kappa, y_{1:\kappa}, x_{1:\kappa}\}$

return ψ

▷ Initialize blank image canvas

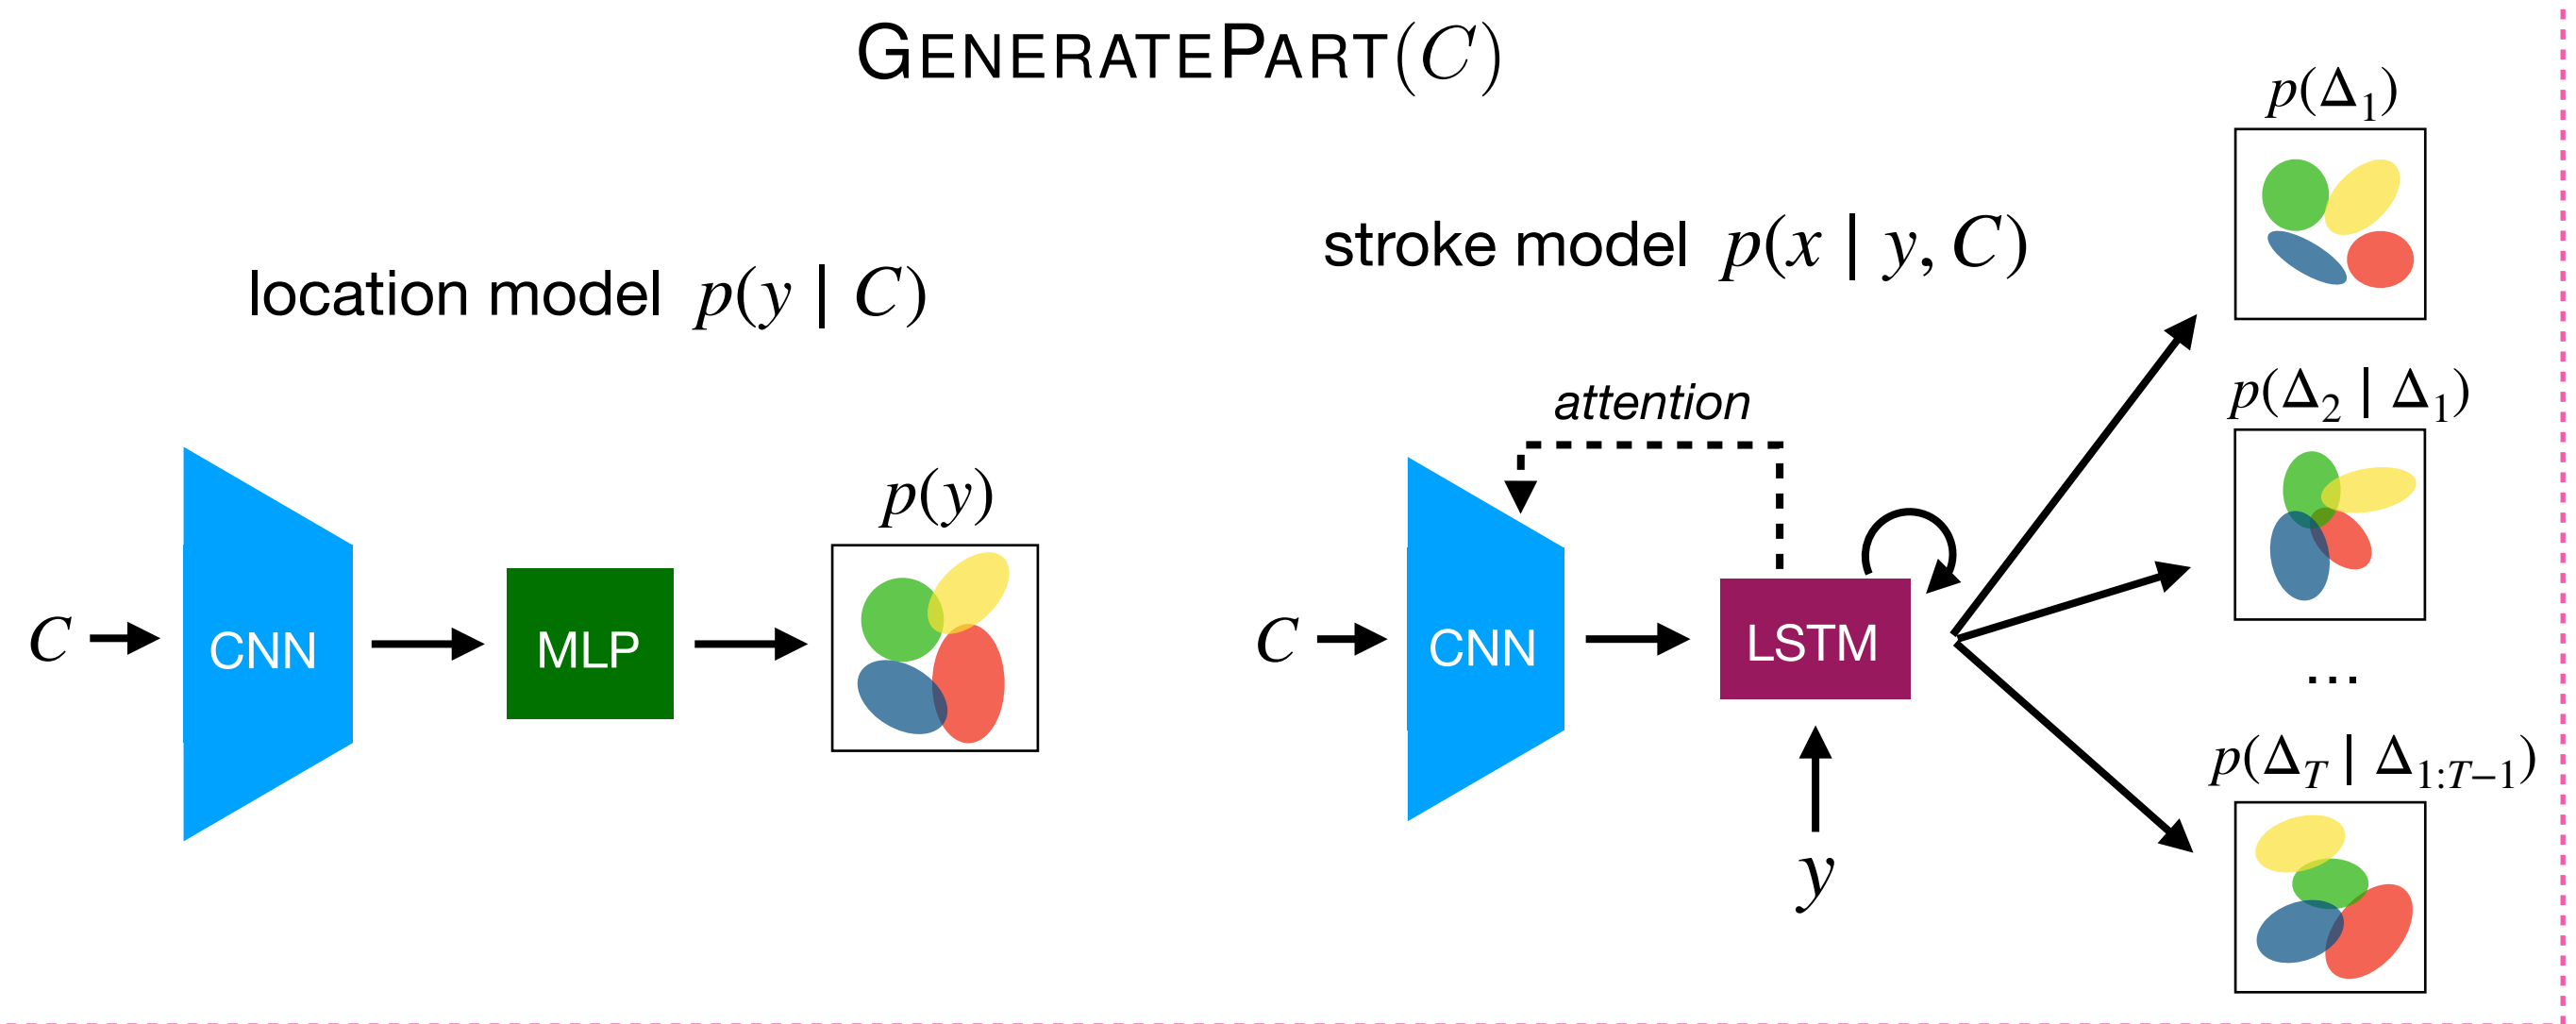
▷ Sample part location & parameters

▷ Render part to image canvas

▷ Sample termination indicator

▷ Terminate sample

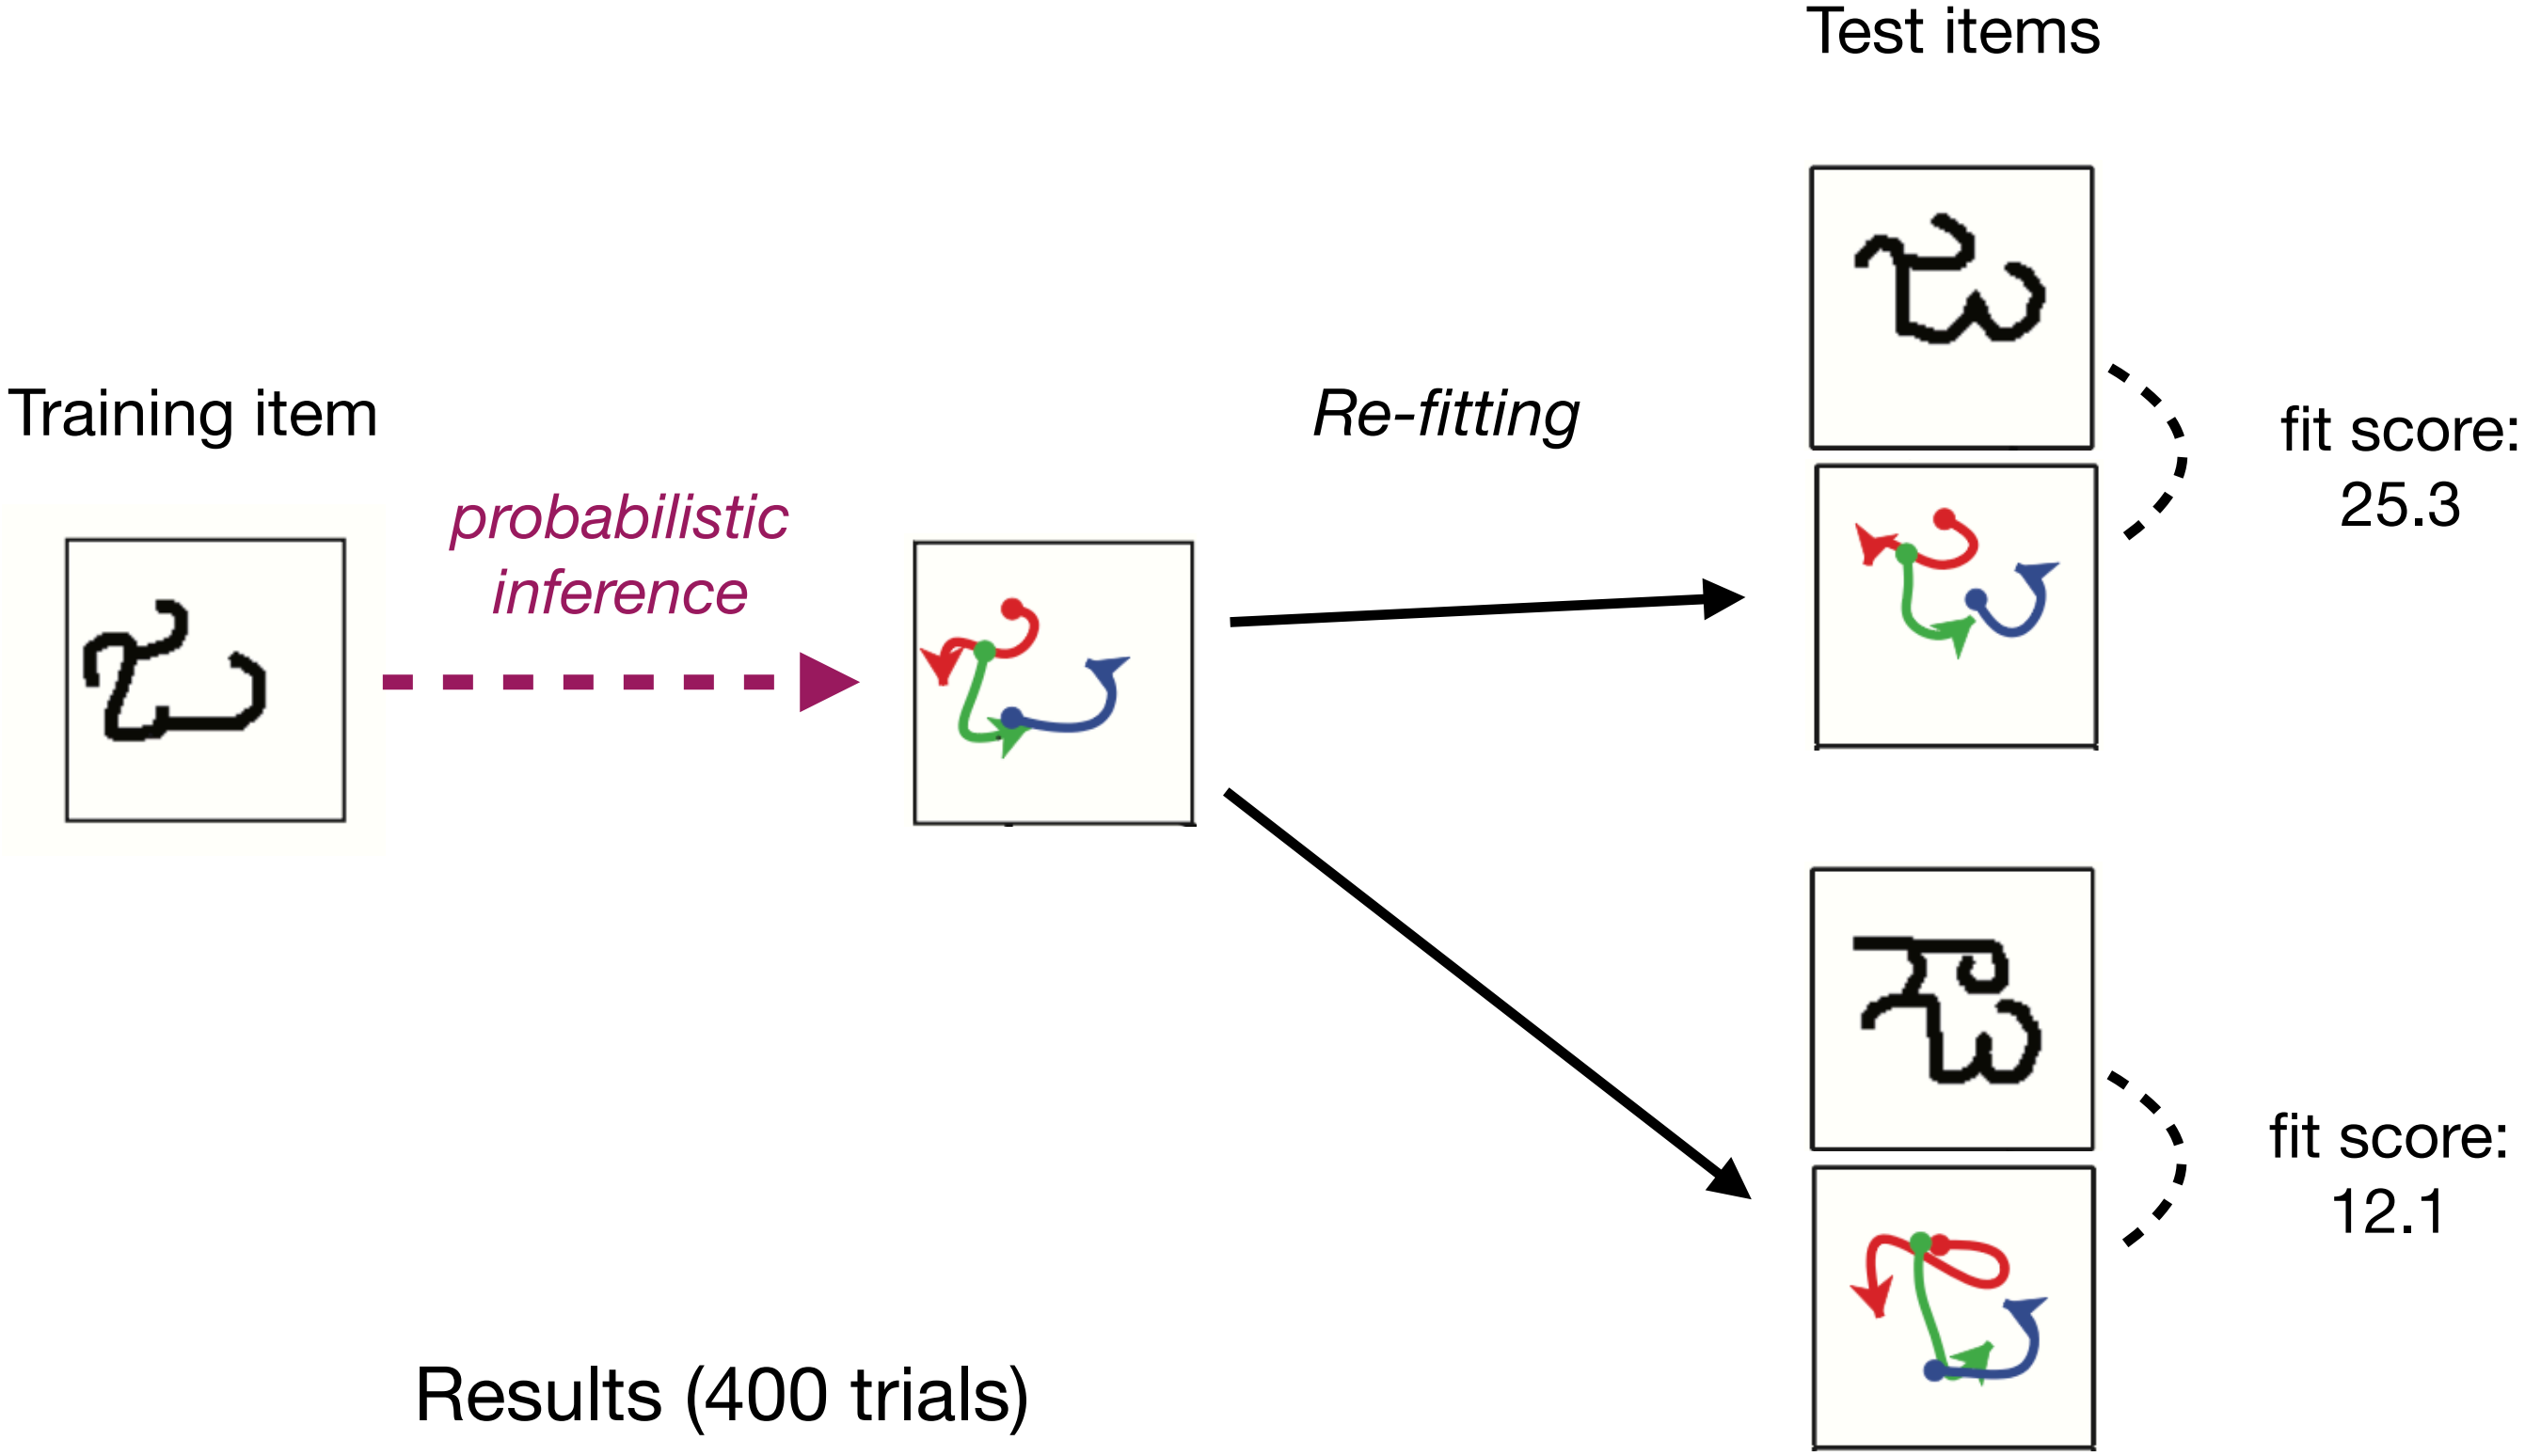
▷ Return concept type



One-Shot Classification



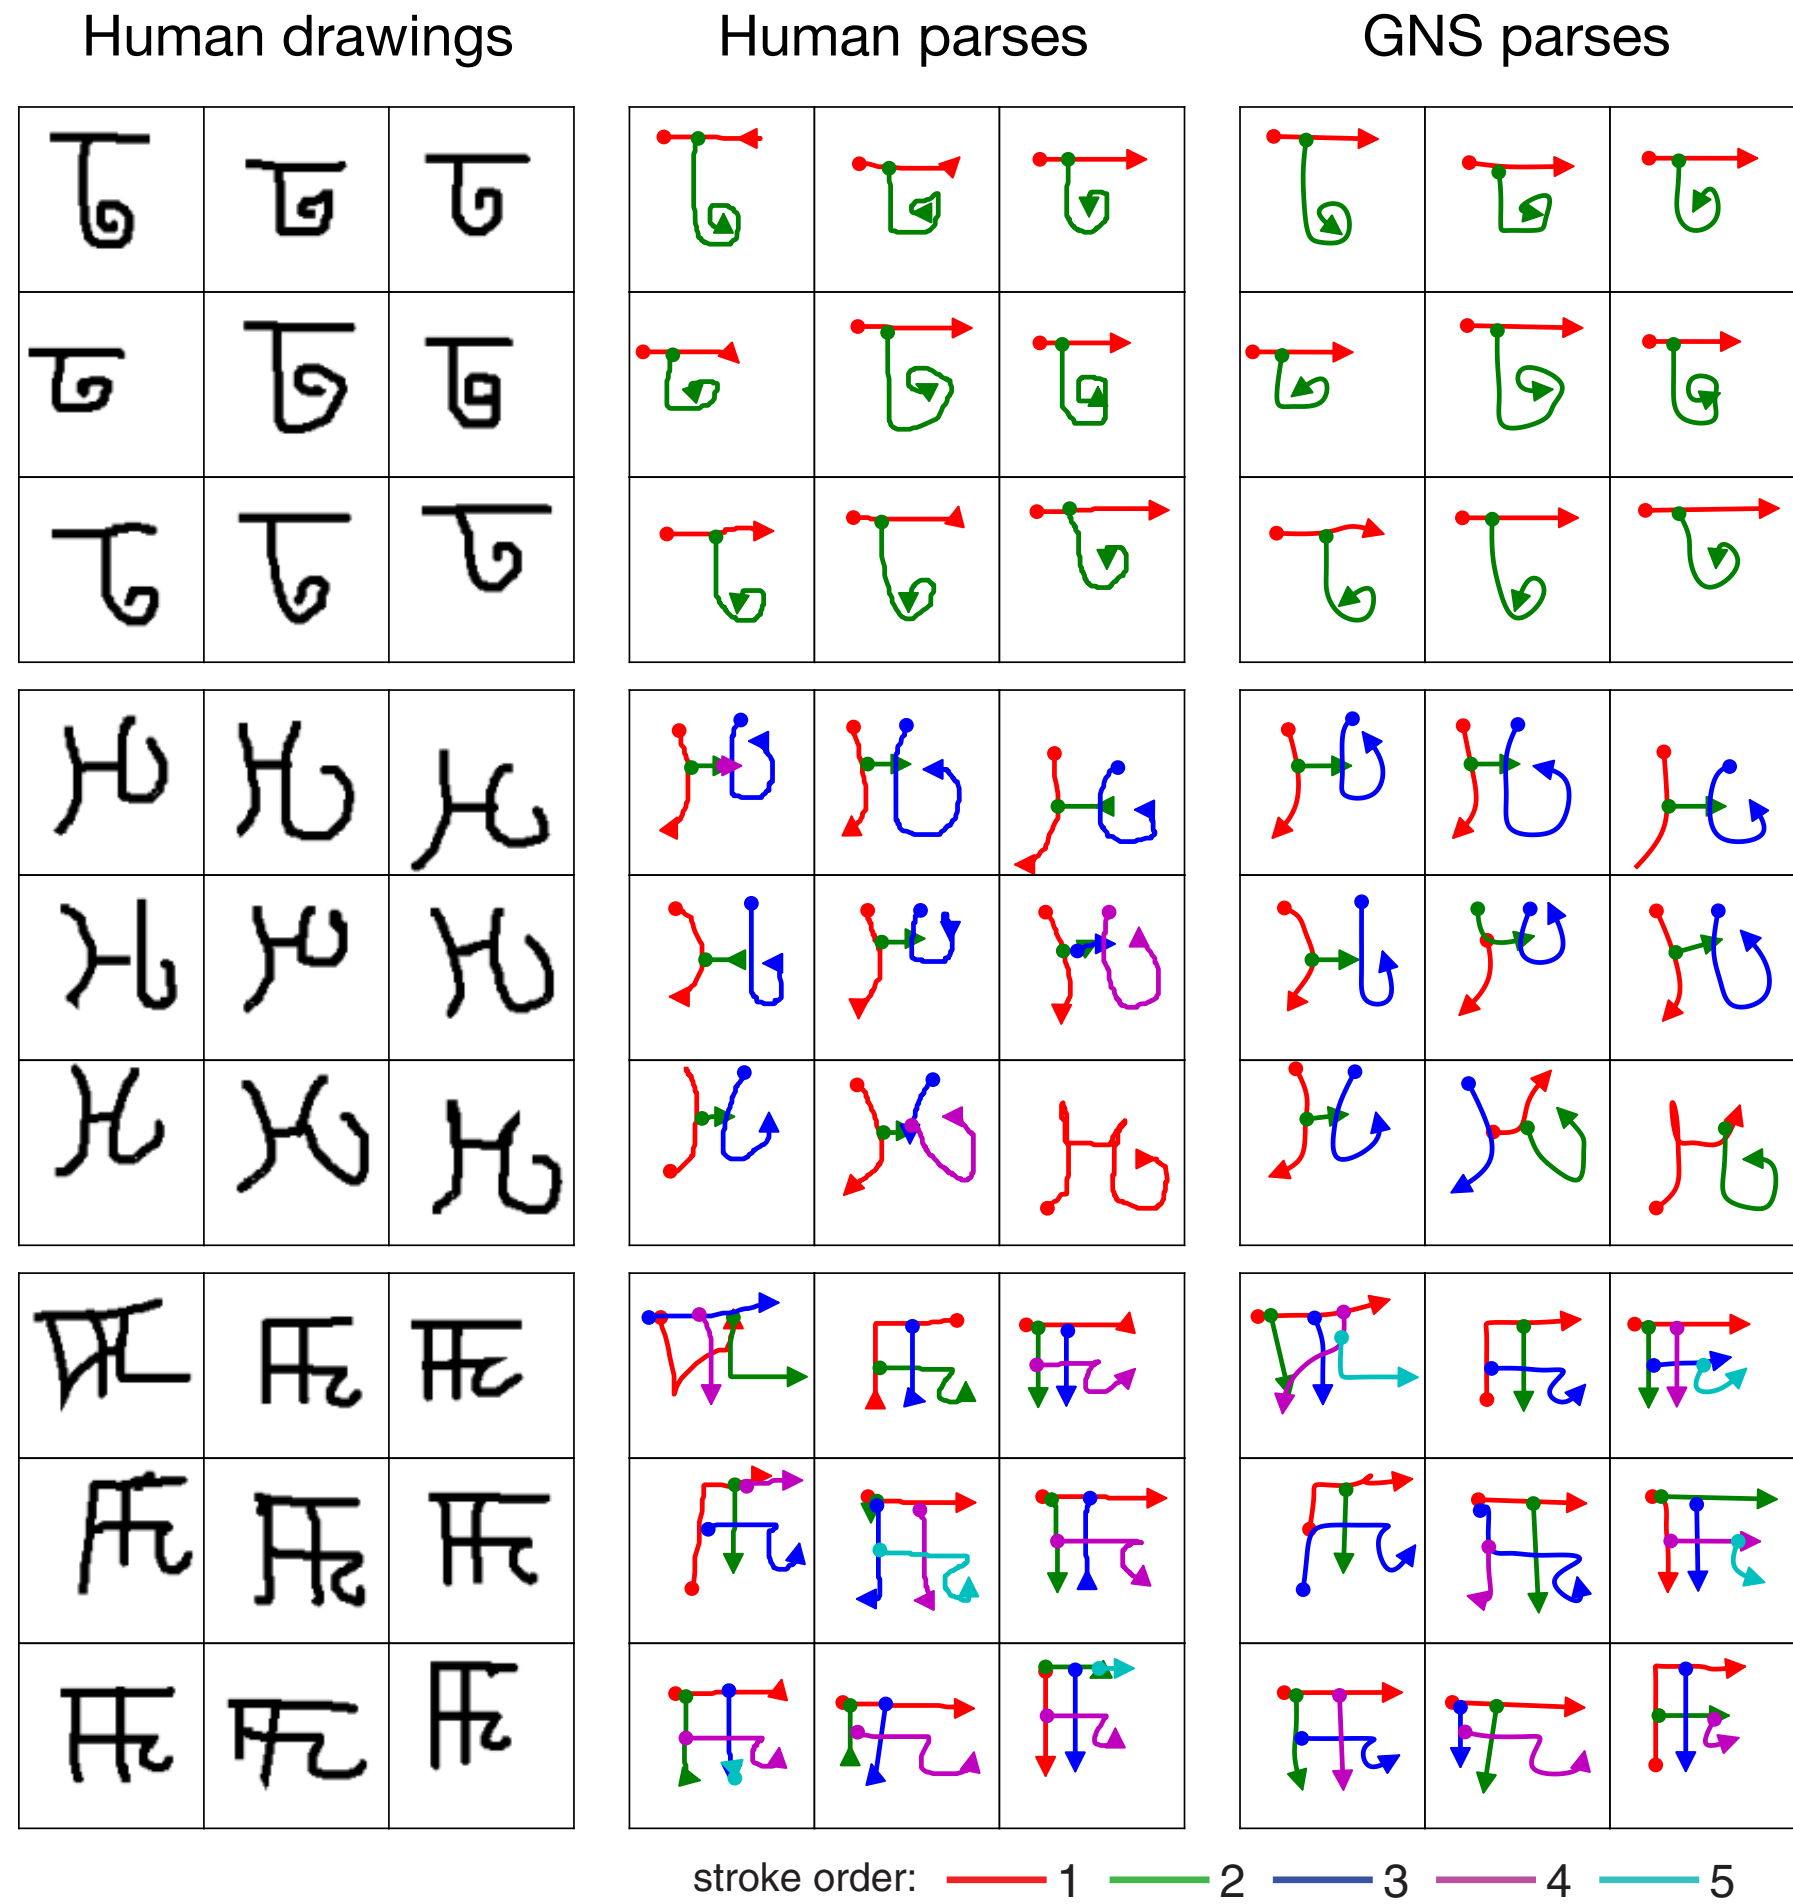
ఓ	ఇ	లు	ఎ	ఔ
అ	బ	గ	ఓ	ఝ
క్ష	ల	ణ	త	ద
న	య	ల	రా	ళ




Results (400 trials)

	Accuracy
GNS	94.3%
Humans	95.5%

Parsing



Generating new exemplars

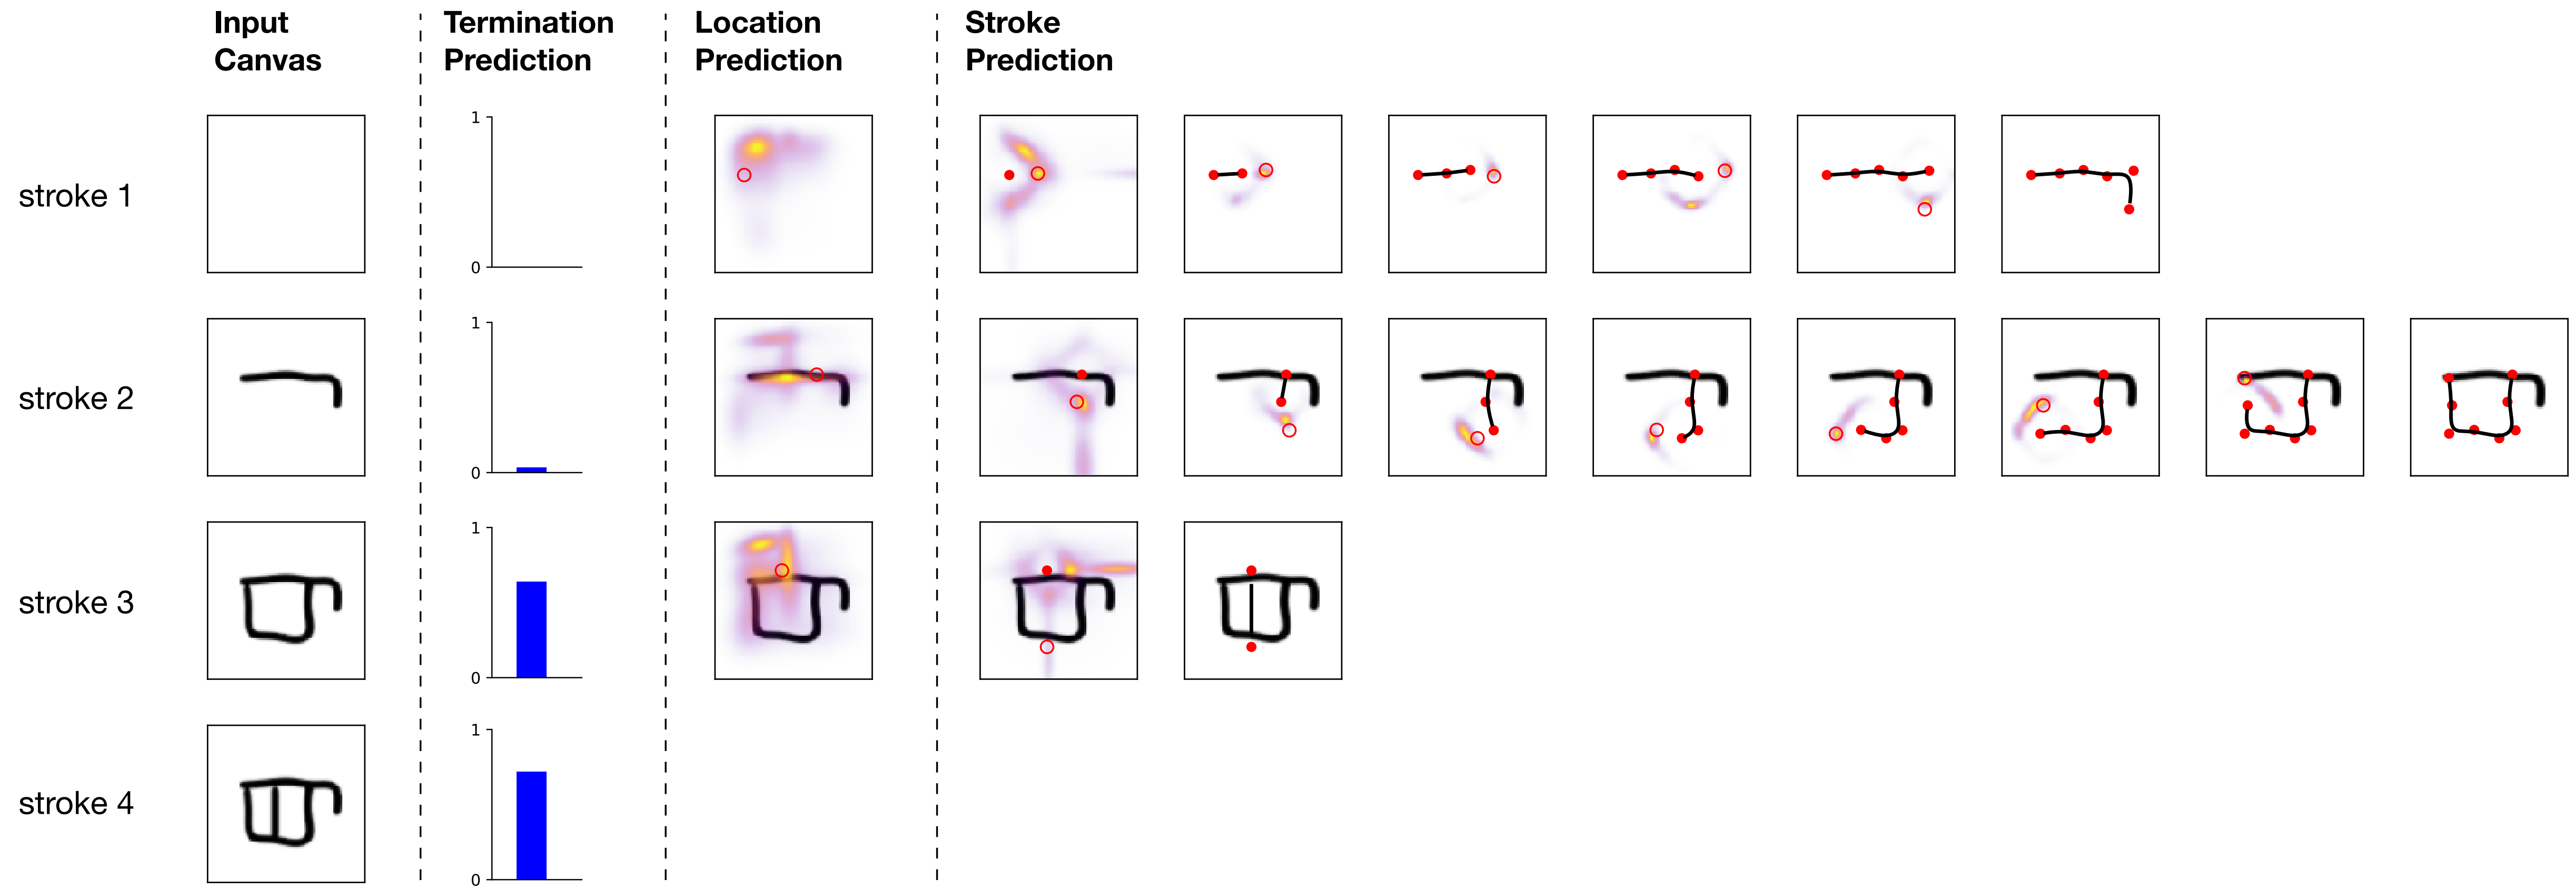
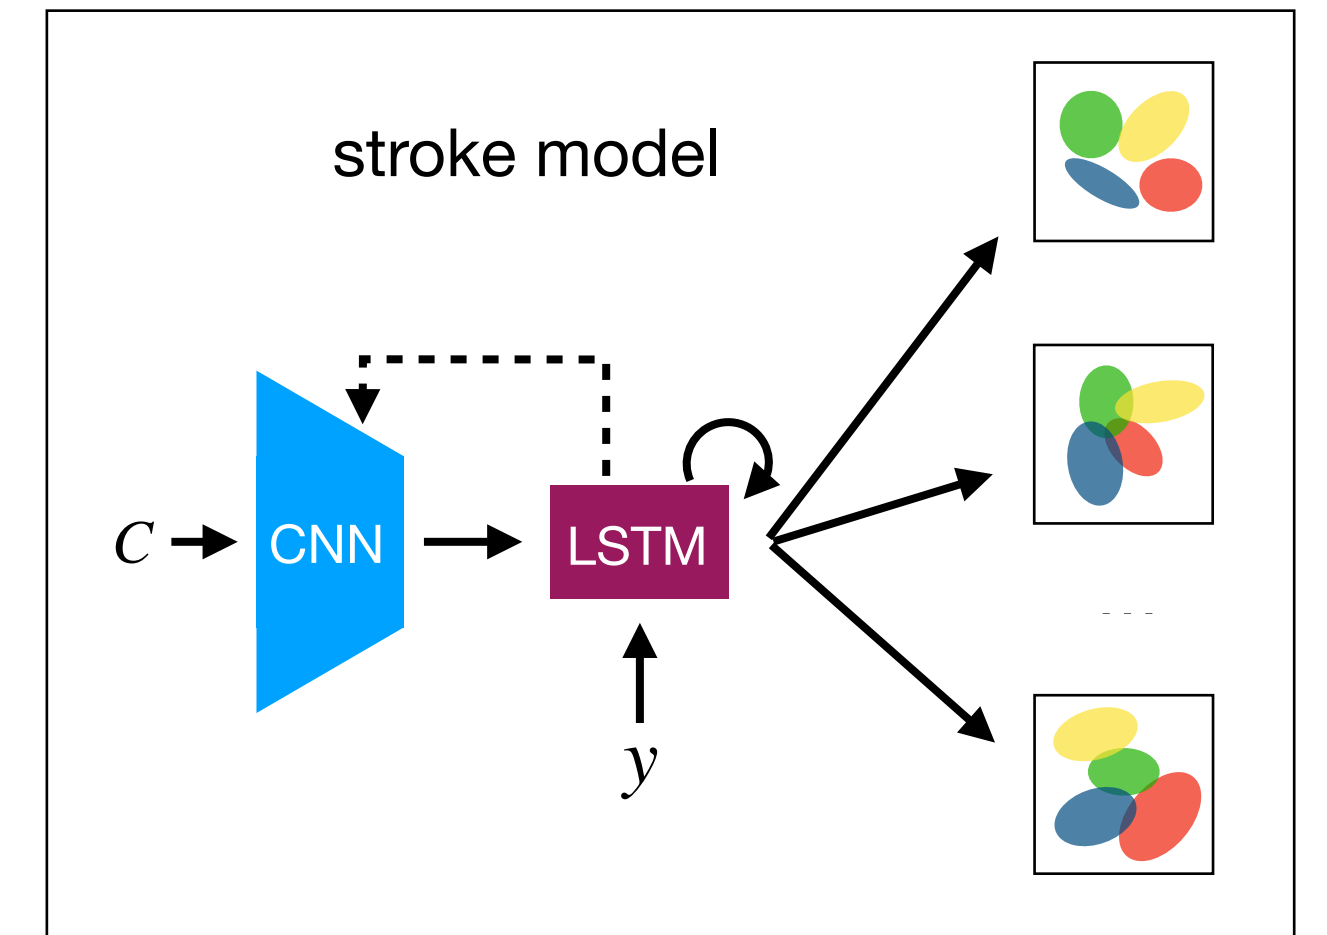
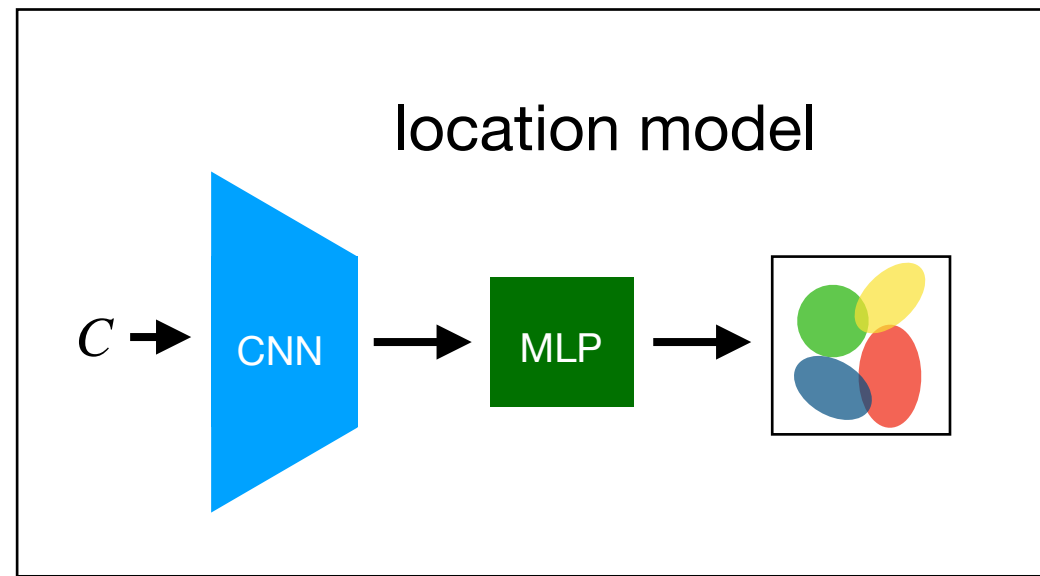


Thanks


For more information, come to my poster!

Extras

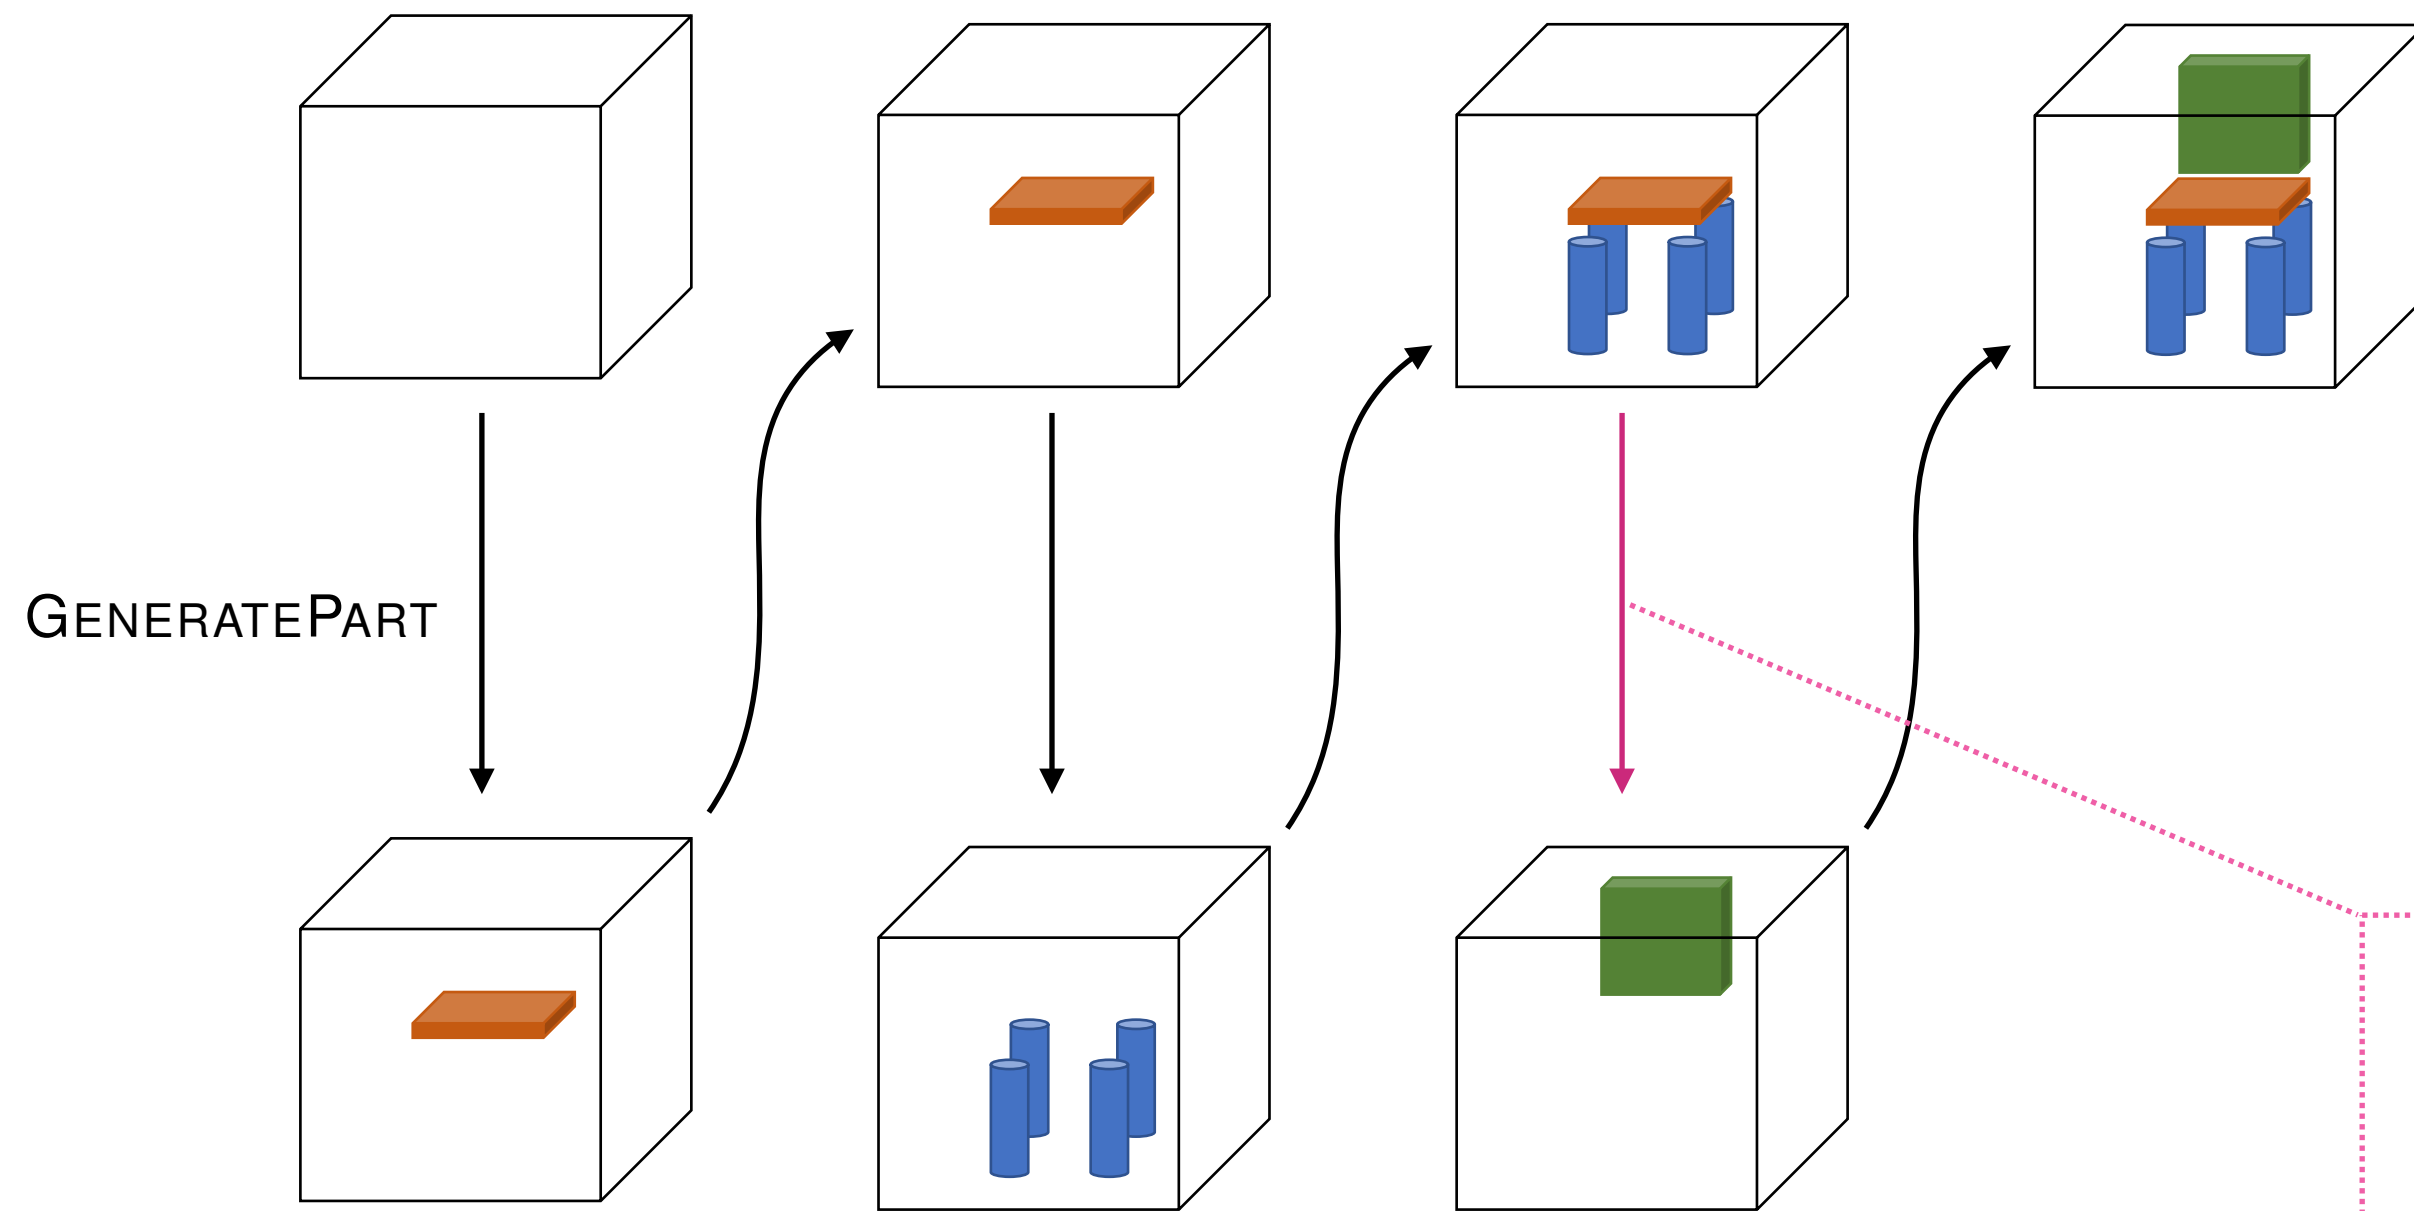
Type Prior



3D objects (ShapeNet)



GNS model of 3D objects



procedure GENERATE_TOKEN

```

C ← 0
while true do
  xi ← GENERATEPART(C)
  C ← fupdate(xi, C)
  vi ~ p(v | C)
  if vi then
    break
θ ← {κ, x1:κ}
return ψ
  
```

- ▷ Initialize blank 3d canvas
- ▷ Sample part from neural net
 - ▷ Update 3d canvas
- ▷ Sample termination indicator
 - ▷ Terminate sample
- ▷ Return concept token

GENERATEPART(C)

