CSC2547: Learning to Search

Intro Lecture
Sept 13, 2019

This week

- Course structure
- Background, motivation, history
- Project guidelines and ideas
- Ungraded quiz

Course Schedule

- Weeks 1 & 2: Intro & Background (by me)
- Weeks 3-10: Paper presentations and tutorials (by you)
- Weeks 11 & 12: Project presentations (by you)

Marks Breakdown

- [15%] Assignment on gradient estimation and tree search
- [15%] 10min class presentations
- [15%] 2-4 page project proposal
- [15%] 5-min project presentations
- [40%] 4-8 page project report and code

Why to take this course

- To learn about this research area, and the relevant tools (e.g. MCTS, Direct Optimization, A* sampling, gradient estimators, REINFORCE, program induction)
- To kick-start a research project
- To learn more about deep learning, reinforcement learning, and discrete optimization
- To improve your presentation skills

Why not to take this course

- To learn about classical AI/search approaches from an expert. See e.g.:
  - Sheila McIlraith:
    - CSC2542: Topics in Knowledge Representation and Reasoning:
      AI Automated Planning, Winter 2019
  - Fahiem Bacchus
    - CSC 2512: Advanced Propositional Reasoning: Winter 2019
- To get help from me with your project / ML application

Focus of Course

- Building adaptive algorithms to search through large, structured, discrete spaces
- Re-using previous or partial solutions on other problems
  - Accelerating classic search algorithms
- Bringing a large-scale continuous optimization perspective to classic AI problems
  - Understanding limitations of relaxation-based approaches
  - Understanding scope and limitations of Monte Carlo Tree Search

Why this topic now?

- Major progress in optimizing large, pure-continuous models. “Success is guaranteed”.
    - Hitting computational bottlenecks due to soft attention that can be address in principle by hard attention
- Interpretability + compactness of discrete representations
- Applications: Optimizing molecules, finding programs, planning, active learning

Why this topic now?

- Eric Langlois working on generalizations of MCTS, need to know current literature.
- Will Saunders working at Ought, raises practical issues of formalizing nested task decomposition.
- Made progress last time on learning with fixed-sized discrete variables (RELAX), got stuck on structured discrete objects like phylogenetic trees

Why this topic now?

- Recent progress, e.g. AlphaZero, Planning chemical synthesis, direct policy gradients

- Existing search strategies are mostly simple and barely adaptive. E.g. reinforce, evolutionary methods, search heuristics

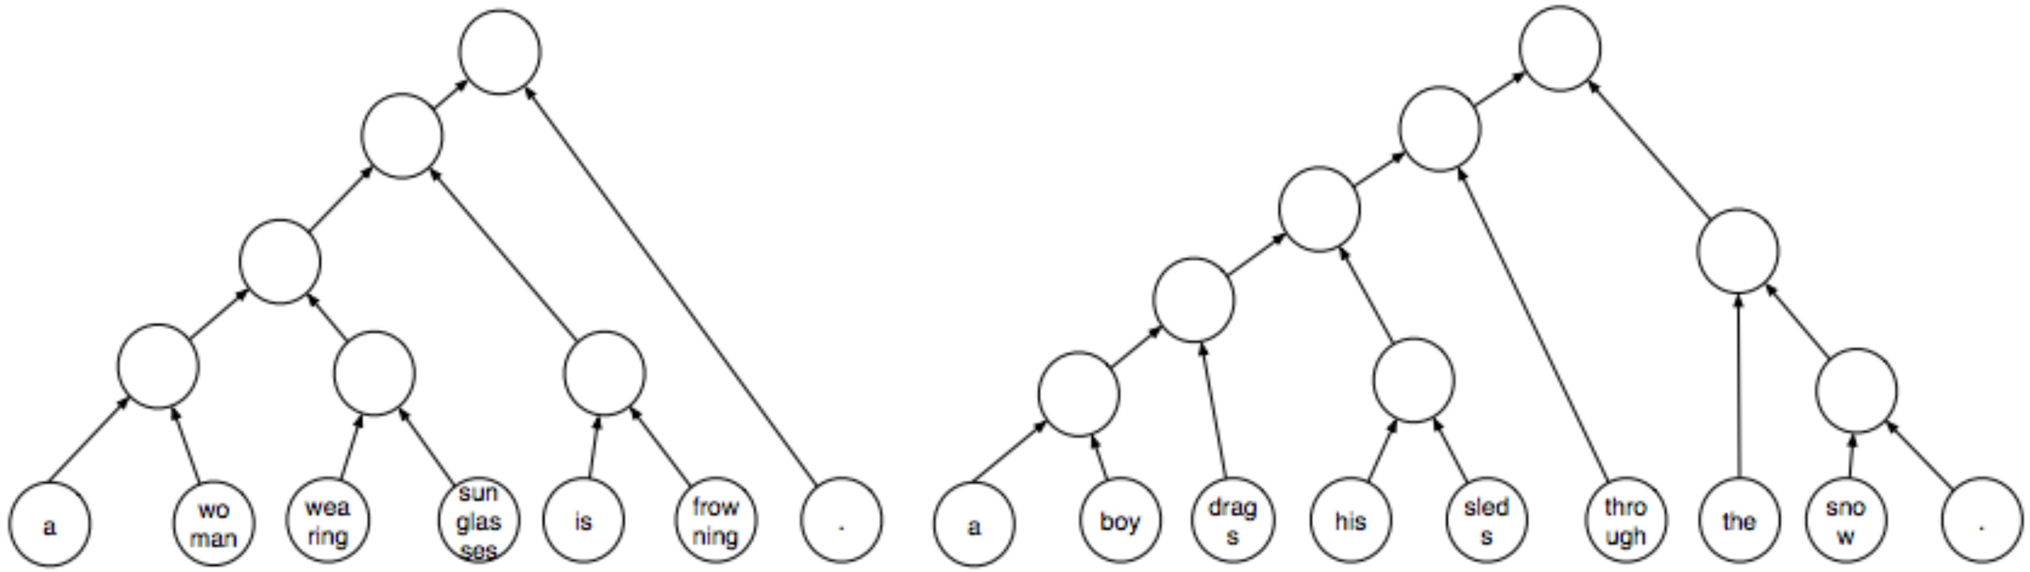

Figure 2: Examples of tree structures learned by our model which show that the model discovers simple concepts such as noun phrases and verb phrases.

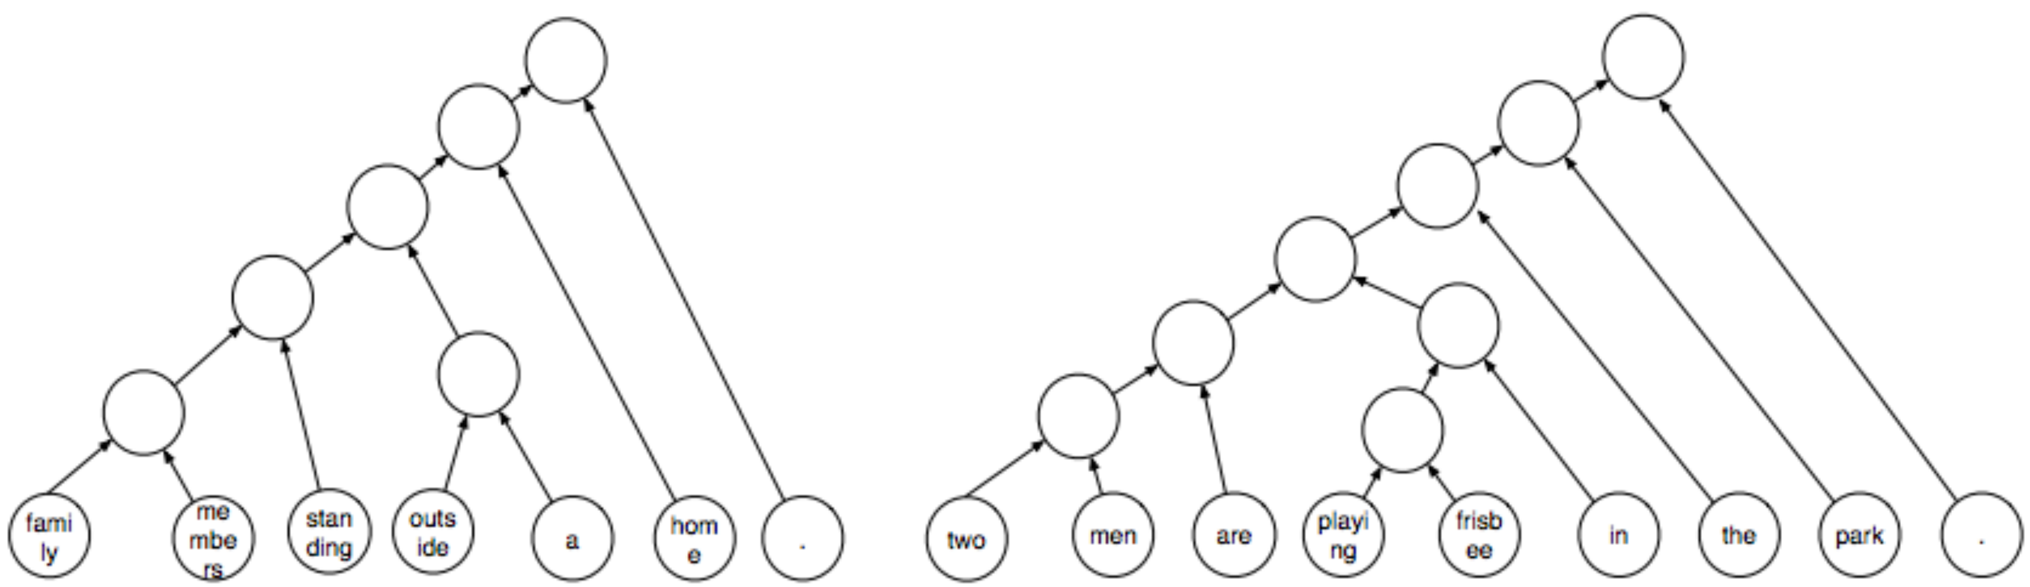

Figure 3: Examples of unconventional tree structures.

Learning to Compose Words into Sentences with Reinforcement Learning
Dani Yogatama, Phil Blunsom, Chris Dyer, Edward Grefenstette, Wang Ling, 2016

```
String s;
BufferedReader br;
FileReader fr;
try {
 fr = new FileReader($String);
 br = new BufferedReader(fr);
 while ((s = br.readLine()) != null) {}
 br.close();
} catch (FileNotFoundException _e) {
  _e.printStackTrace();
} catch (IOException _e) {
  _e.printStackTrace();
}
```

(a)

```
String s;
BufferedReader br;
FileReader fr;
try {
 fr = new FileReader($File);
 br = new BufferedReader(fr);
 while ((s = br.readLine()) != null){}
 br.close();
} catch (FileNotFoundException _e){
} catch (IOException _e){
}
```

(b)

Figure 7: Programs generated in a typical run of BAYOU, given the API method name `readLine` and the type `FileReader`.

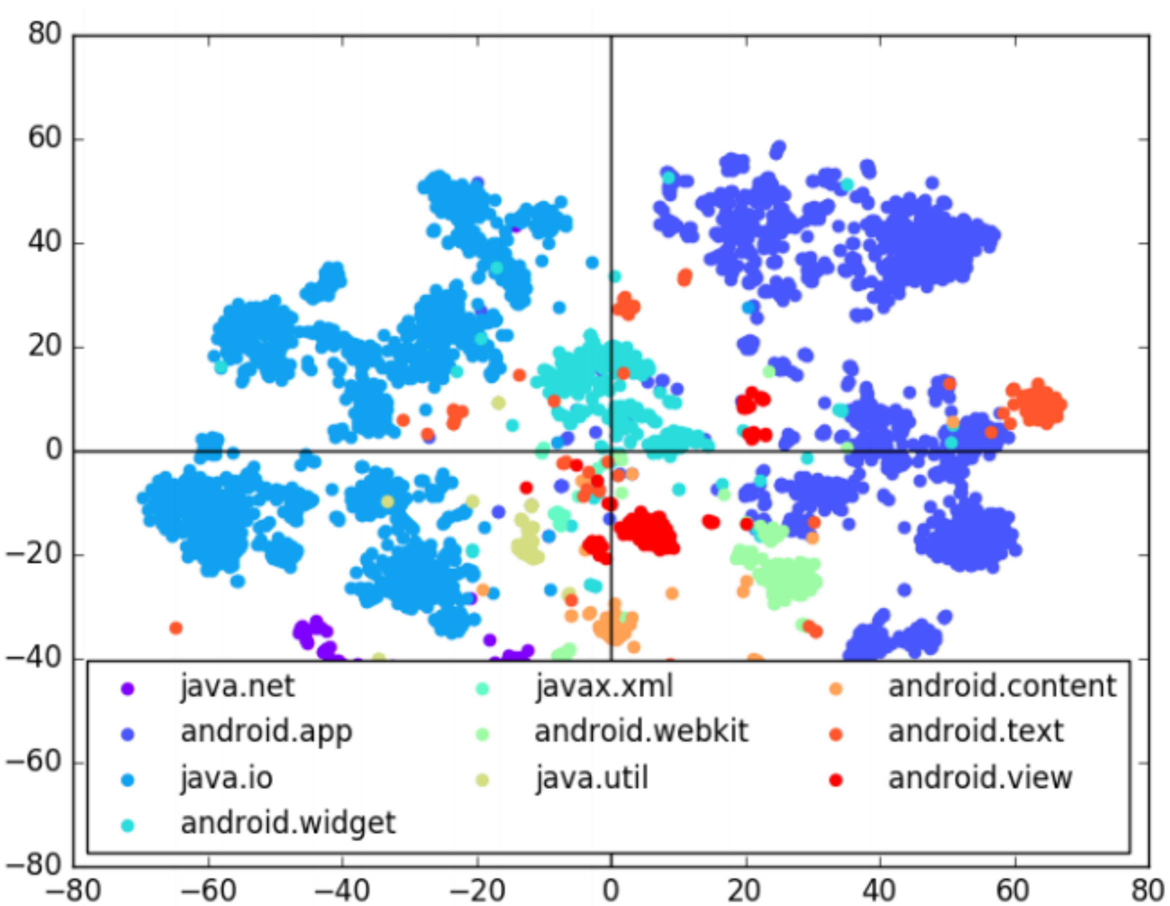

Neural Sketch Learning for Conditional Program Generation, ICLR 2018 submission

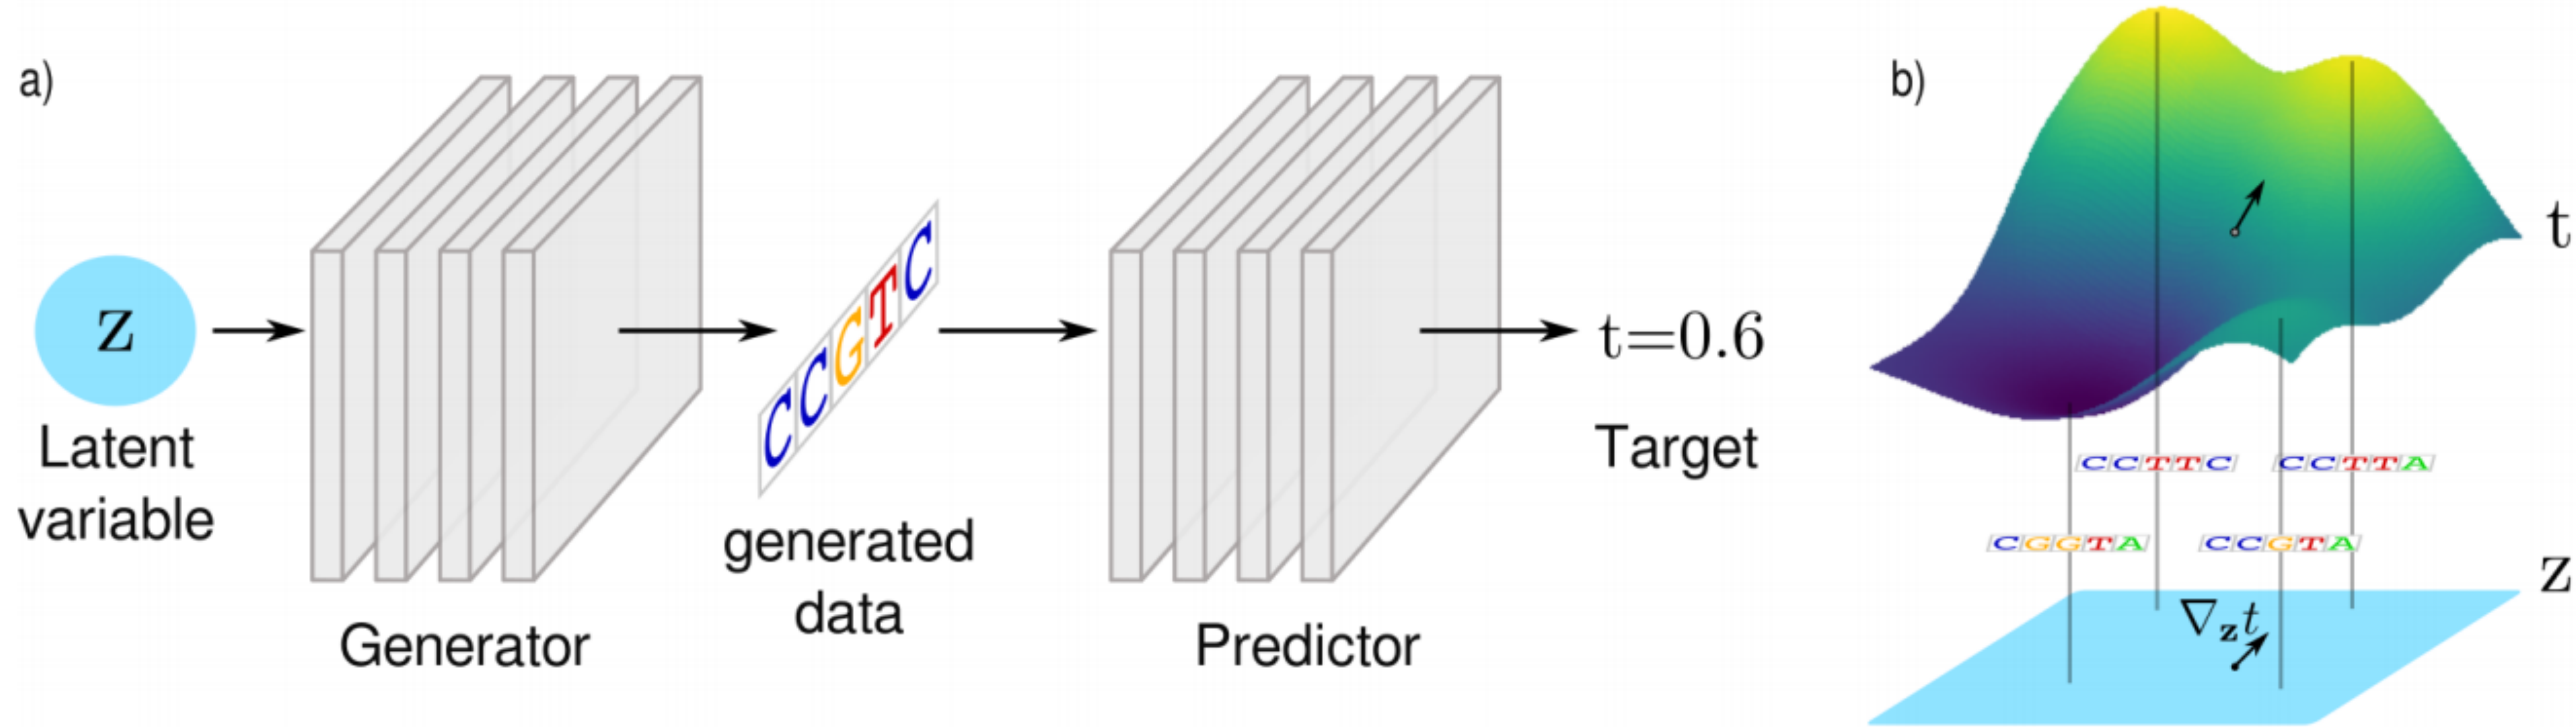

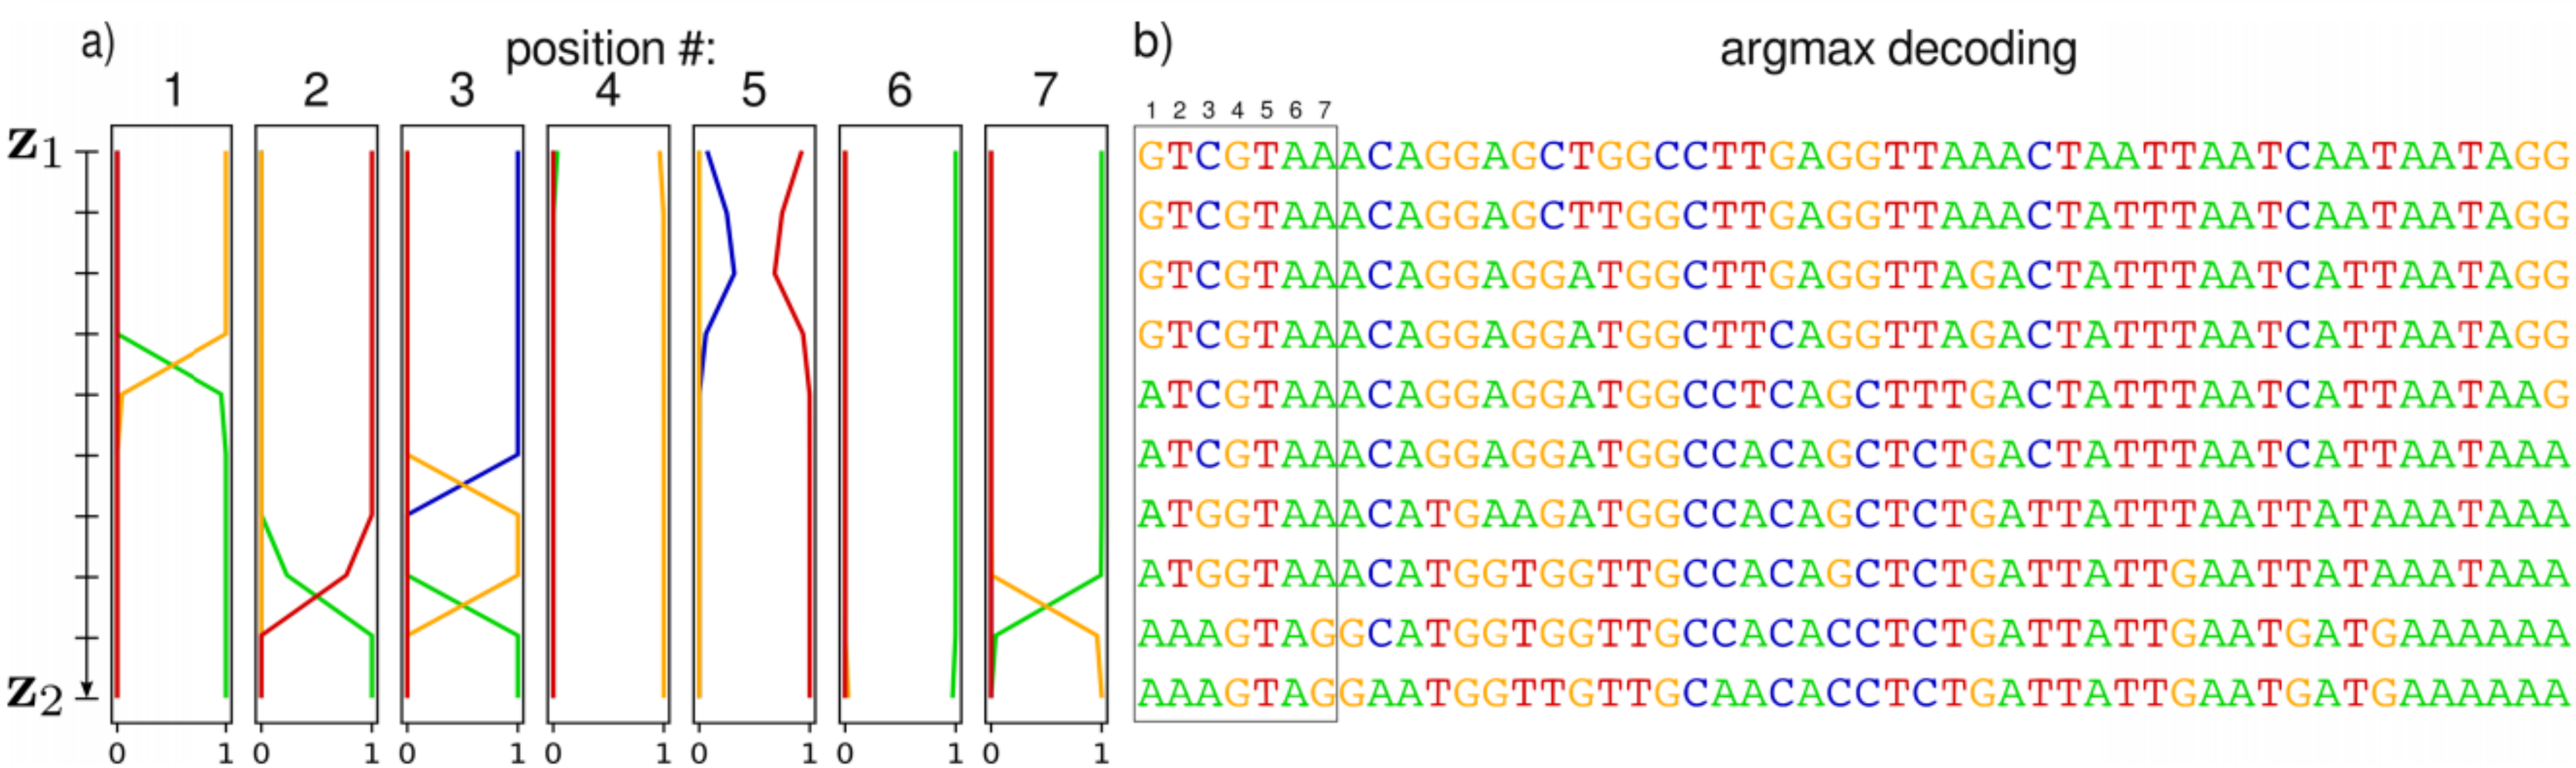

Generating and designing DNA with deep generative models. Killoran, Lee, Delong, Duvenaud, Frey, 2017

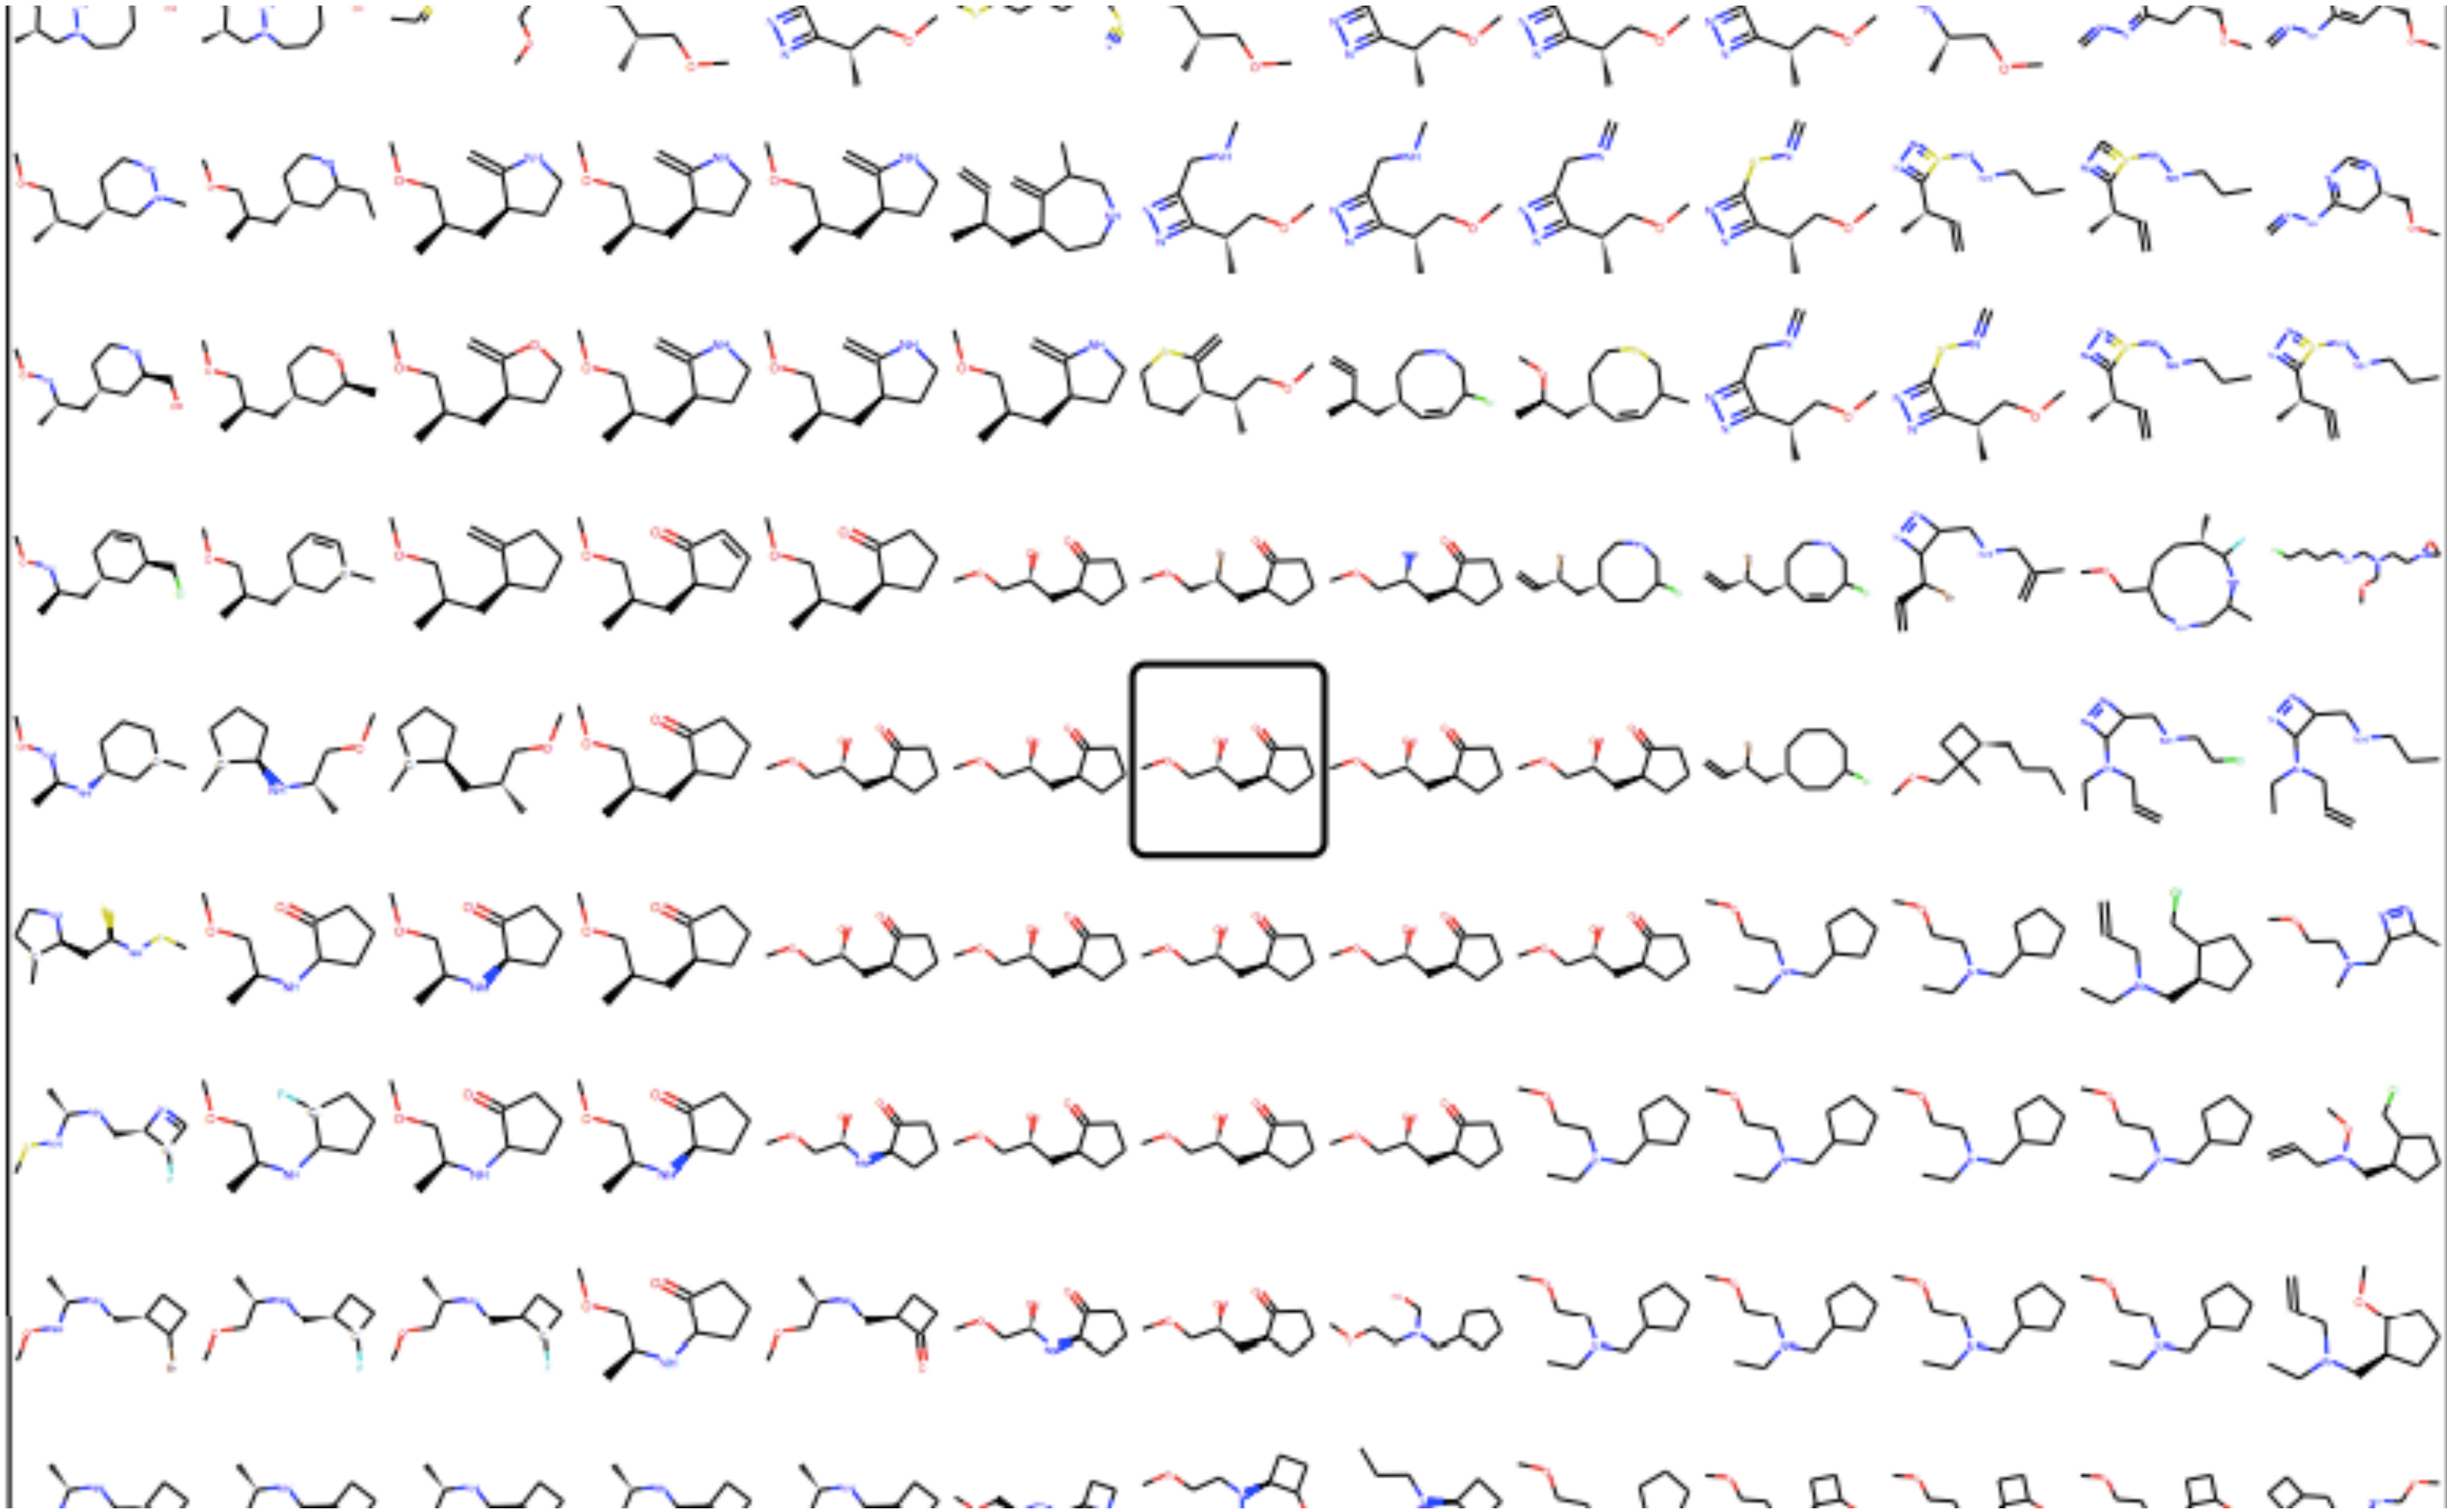

Grammar VAE

Matt Kusner, Brooks Paige, José Miguel Hernández-Lobato

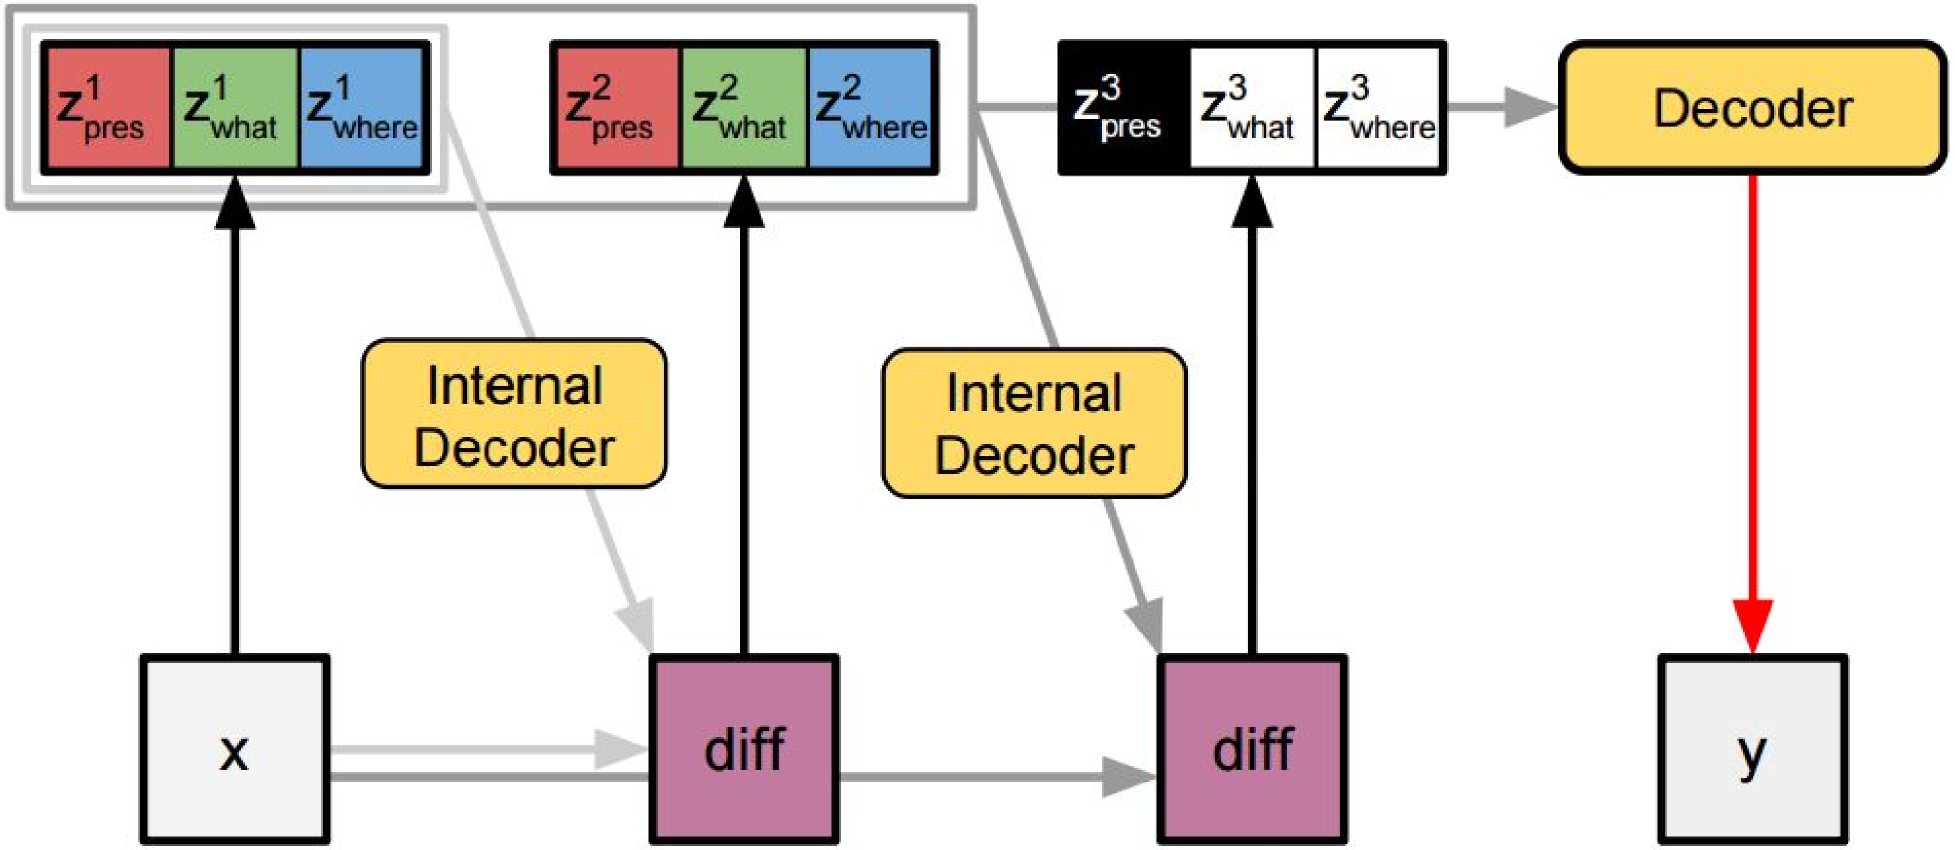

Attend, Infer, Repeat: Fast Scene Understanding with Generative Models

S.M. Eslami,N. Heess, T. Weber, Y. Tassa, D. Szepesvari, K.Kavukcuoglu, G. E. Hinton

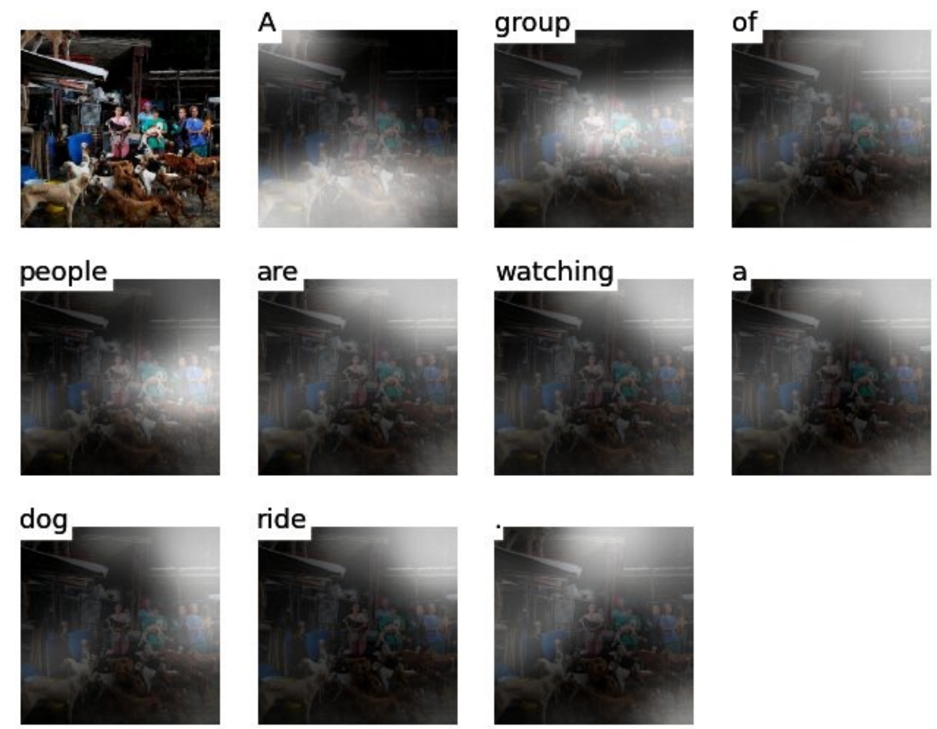

A group of people are watching a dog ride

(Jamie Kyros)

Hard attention models

- Want large or variable-sized memories or ‘scratch pads’
- Soft attention is a good computational substrate, scales linearly O(N) with size of model
- Want O(1) read/write
- This is “hard attention”

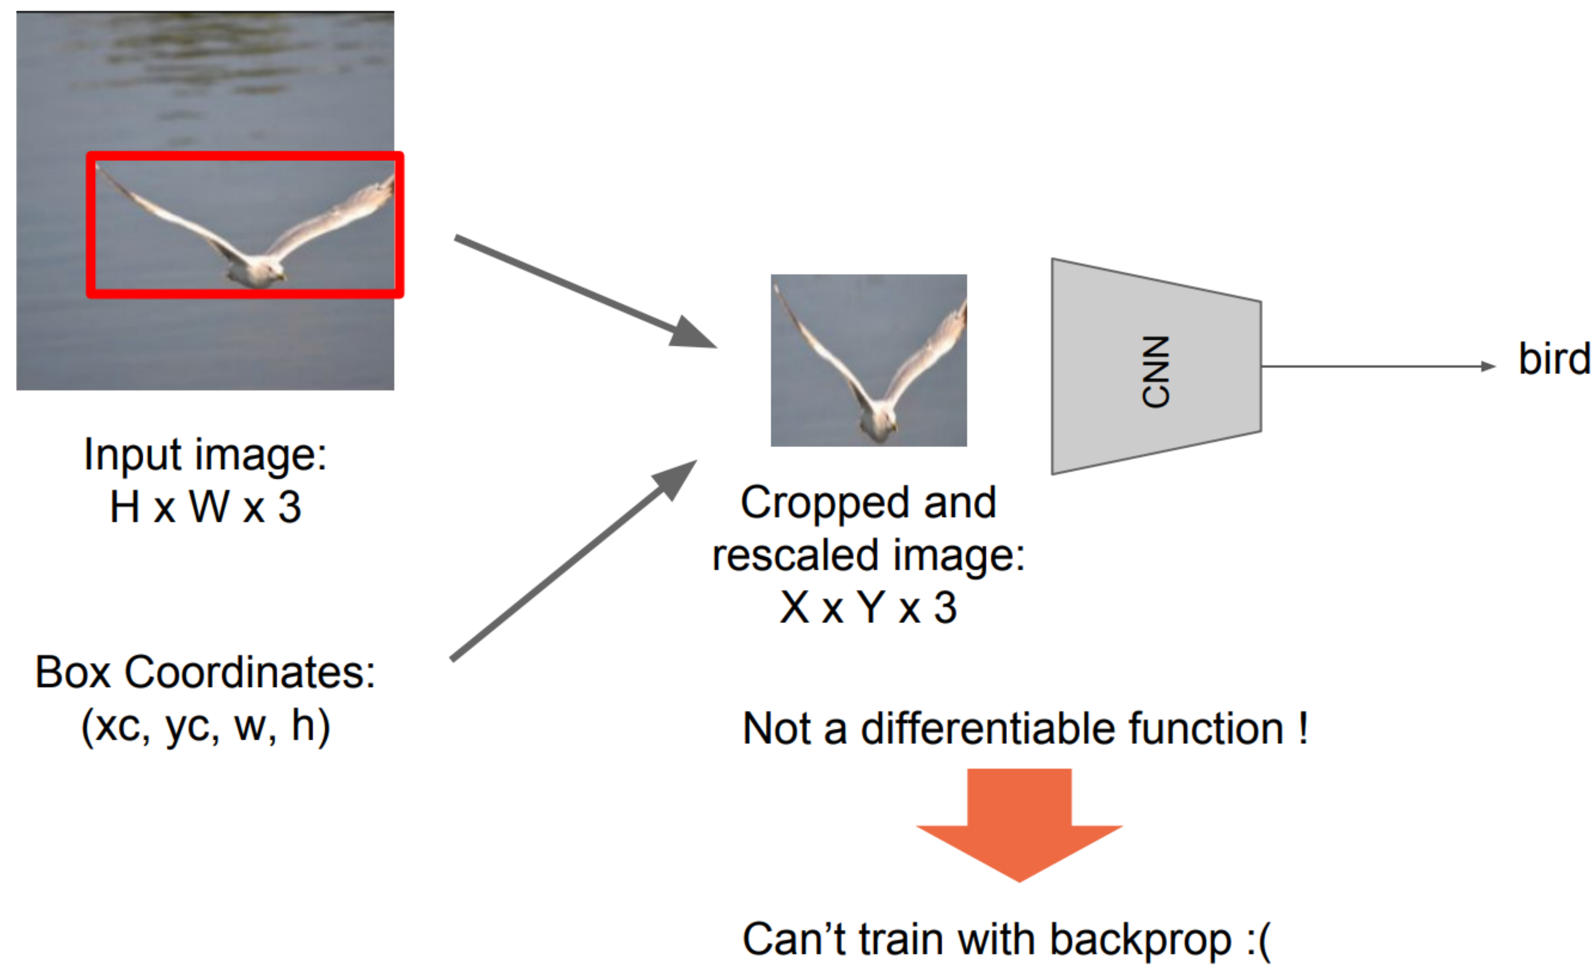

Source: http://imatge-upc.github.io/telecombcn-2016-dlcv/slides/D4L6-attention.pdf

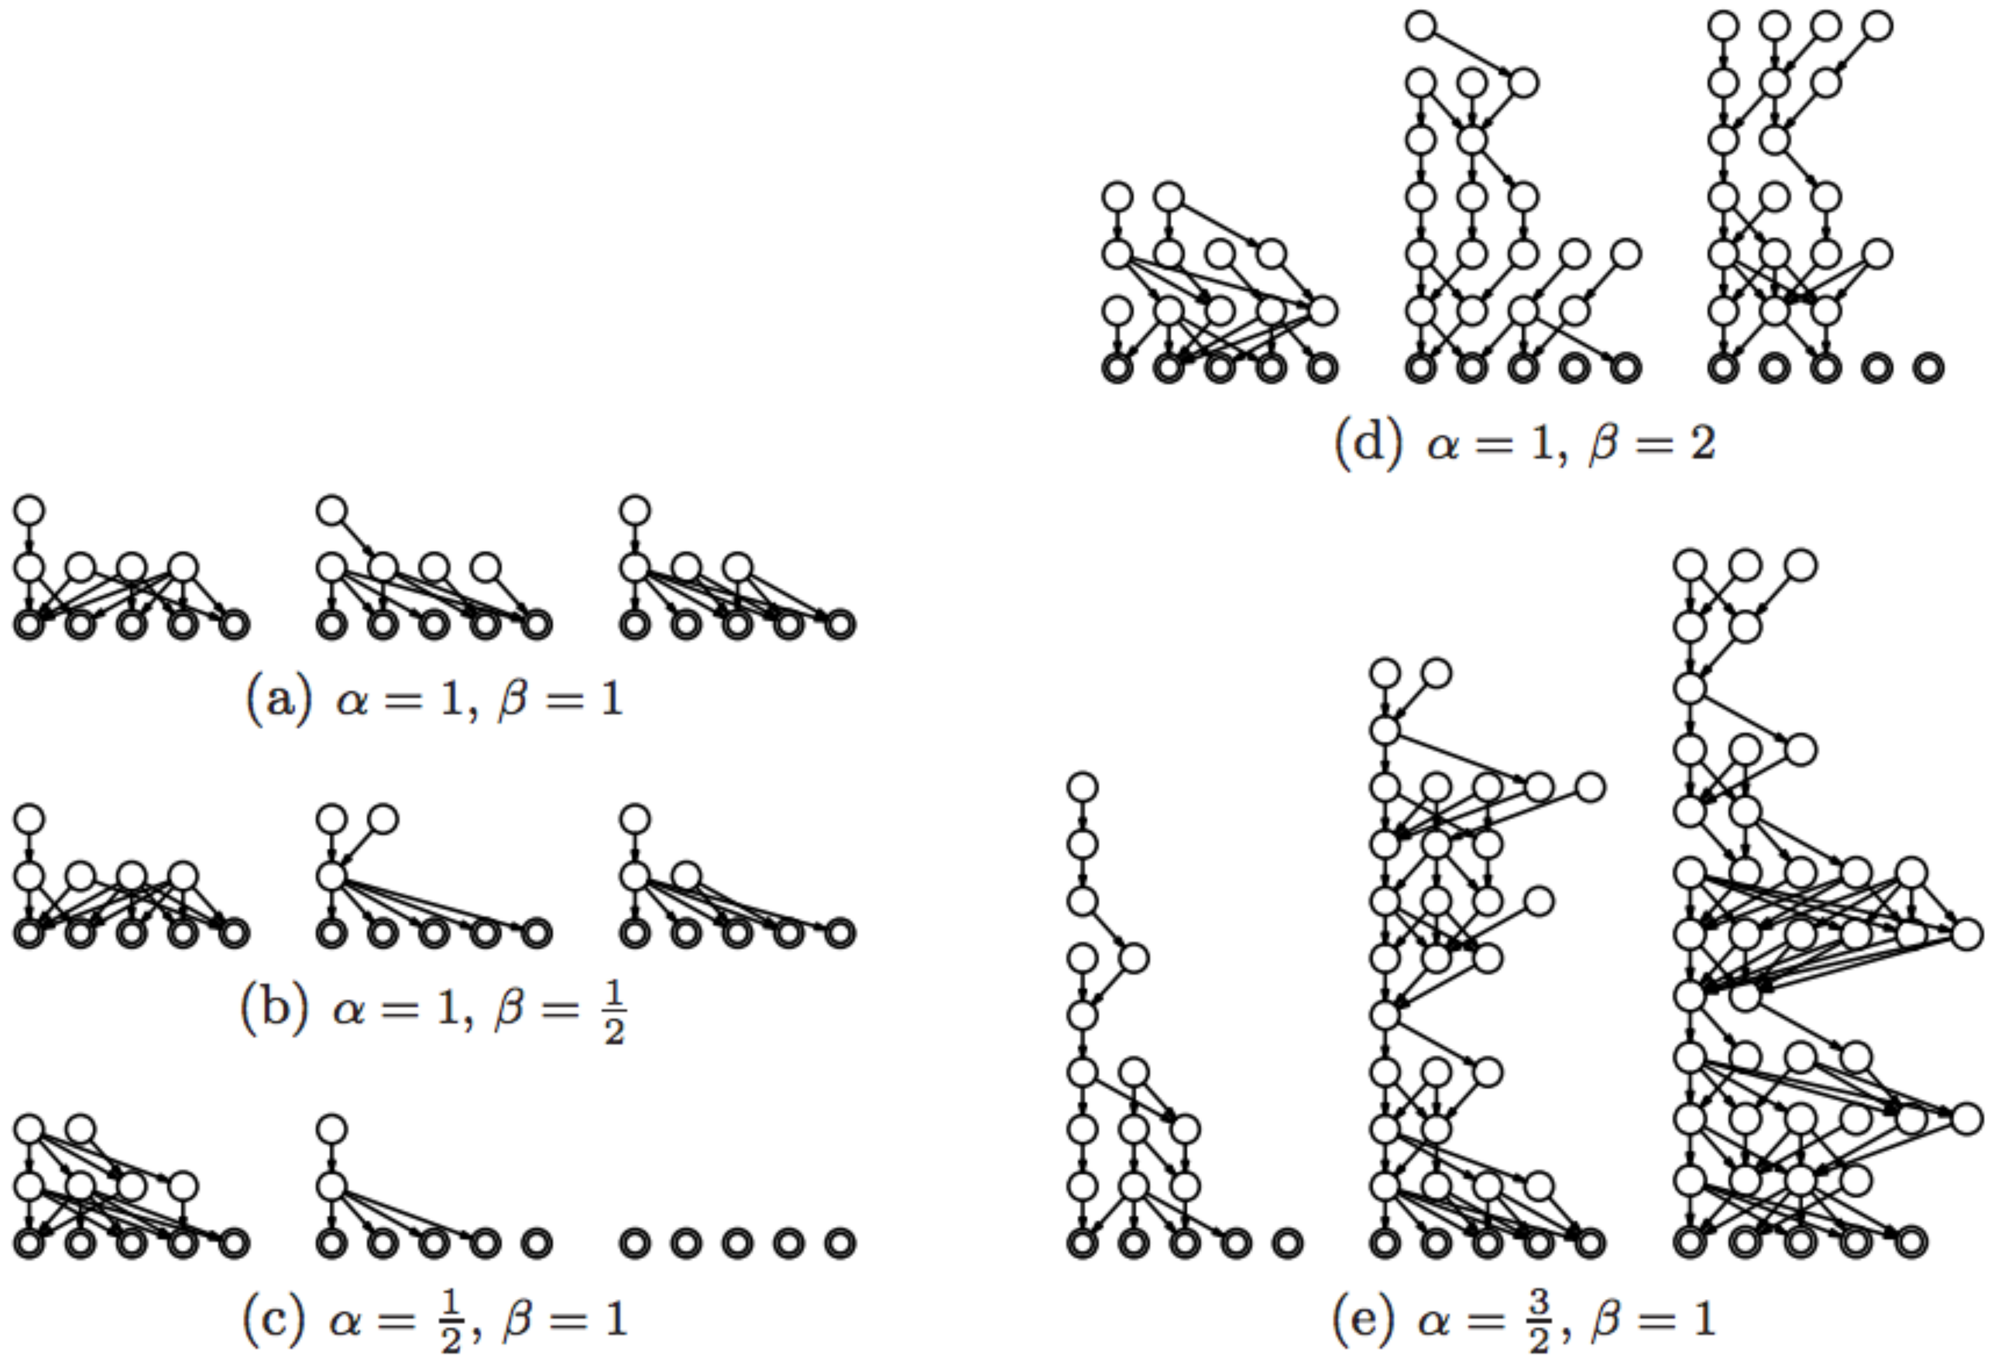

Fig 3: Samples from the CIBP-based prior on network structures, with five visible units.

Learning the Structure of Deep Sparse Graphical Models

Ryan Prescott Adams, Hanna M. Wallach, Zoubin Ghahramani, 2010

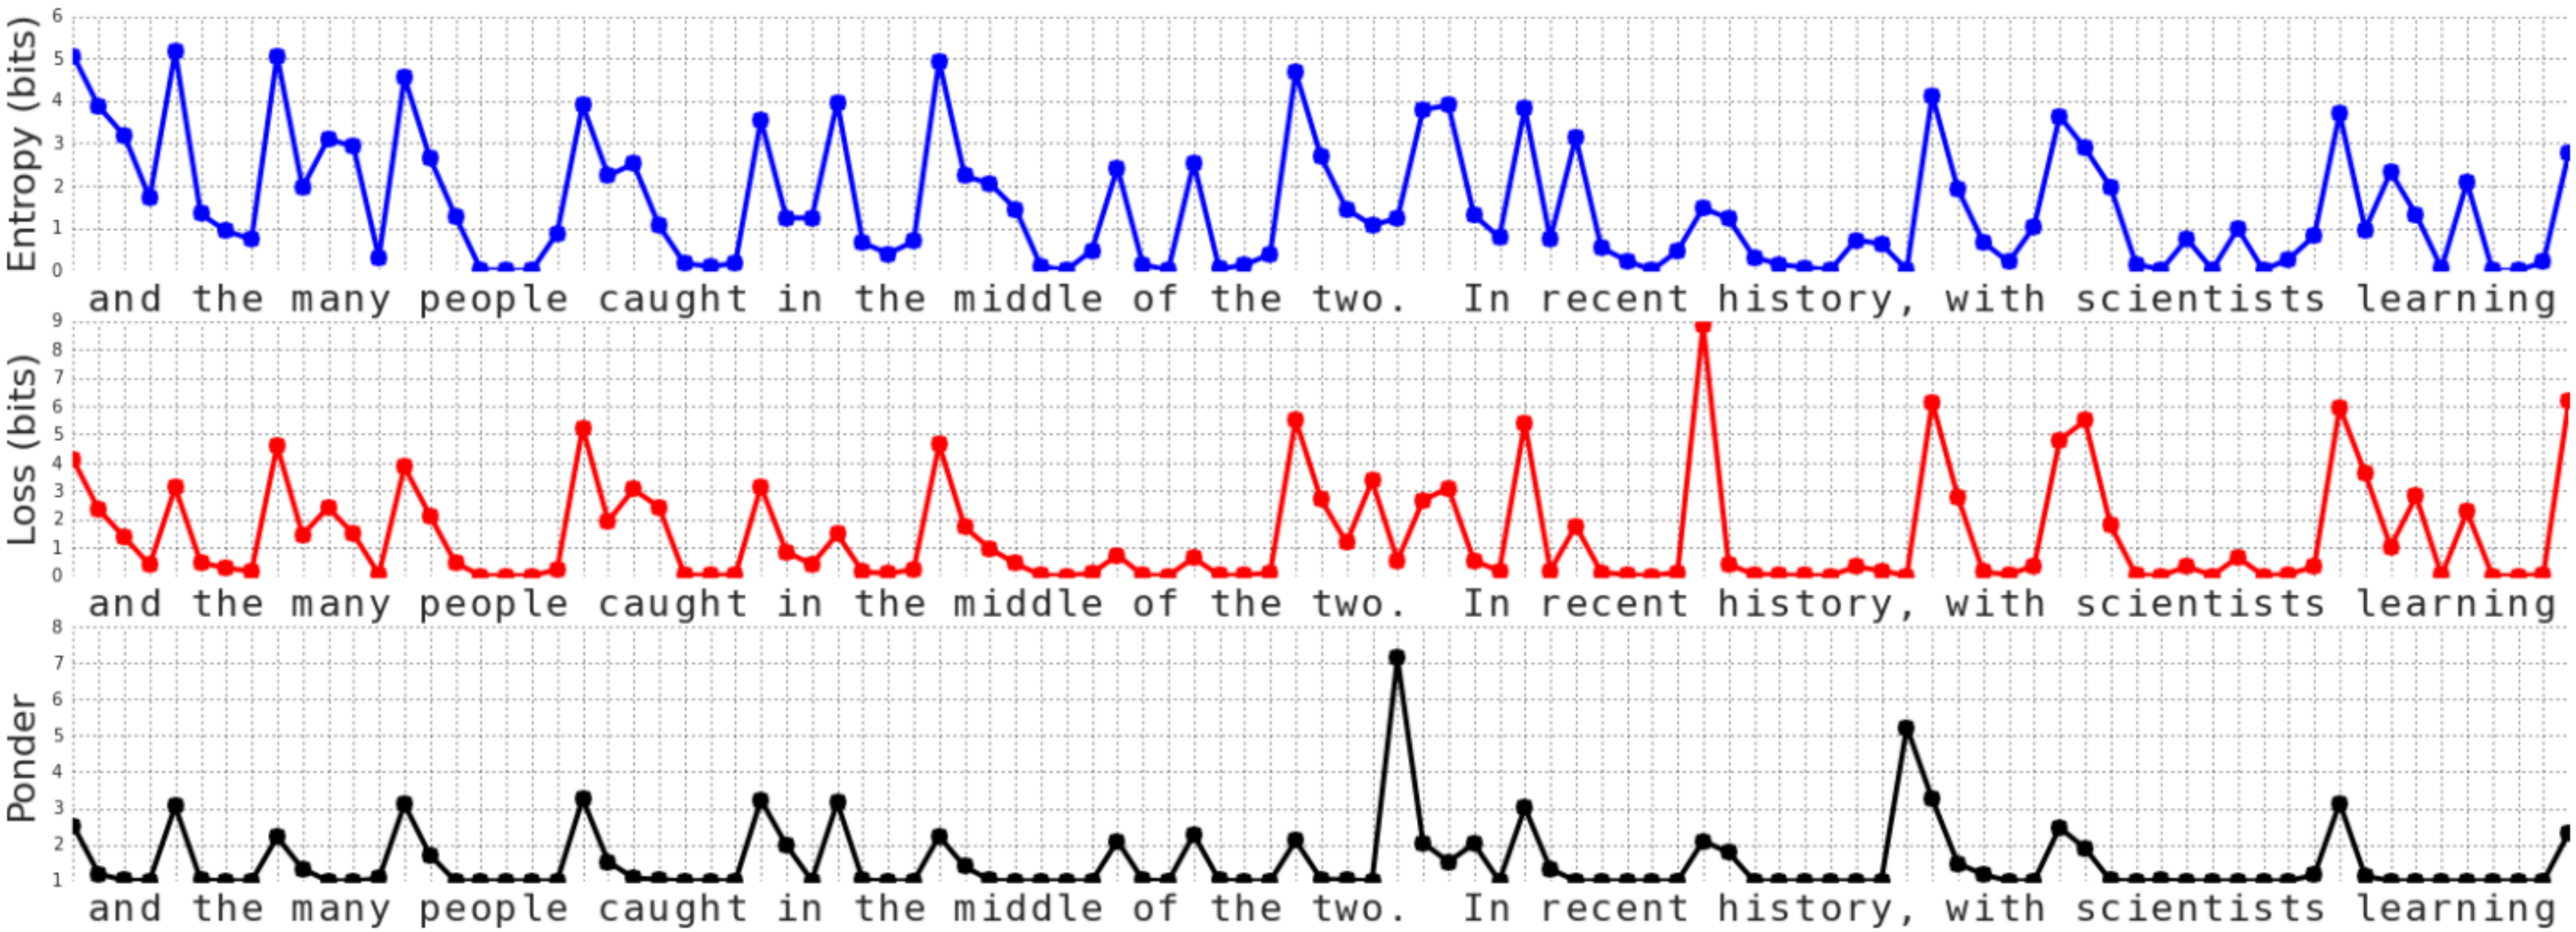

Figure 23: Ponder Time, Prediction loss and Prediction Entropy During a Wikipedia Text Sequence. Plot created using a network trained with $\tau = 6e^{-3}$

Adaptive Computation Time for Recurrent Neural Networks
Alex Graves, 2016

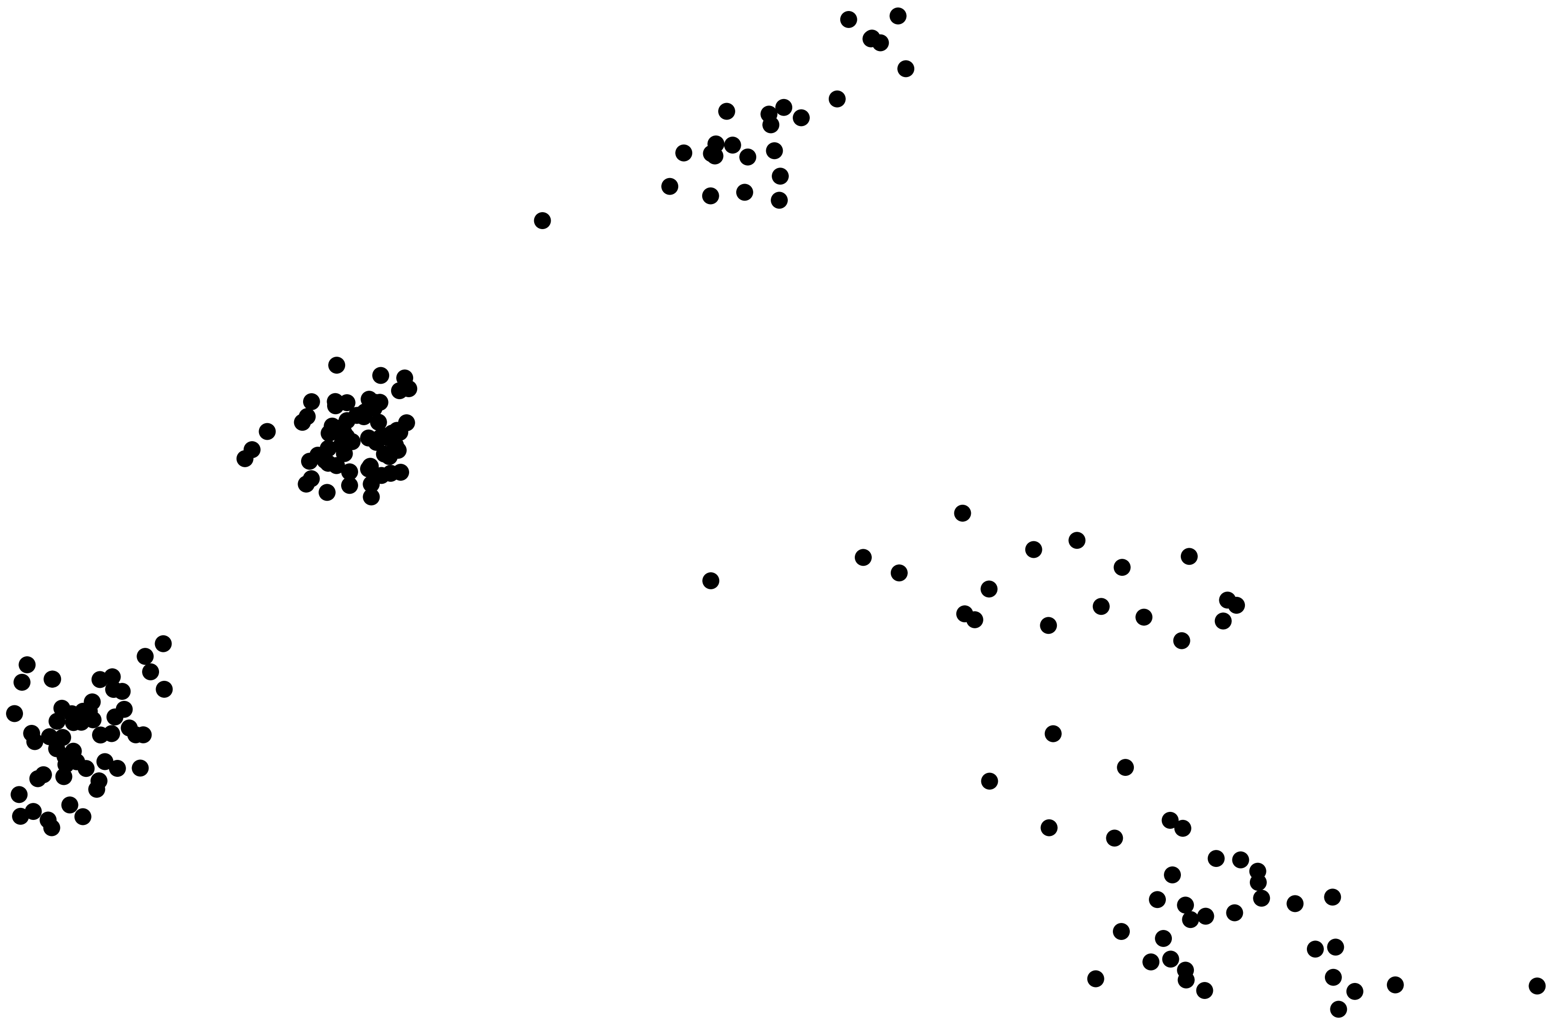

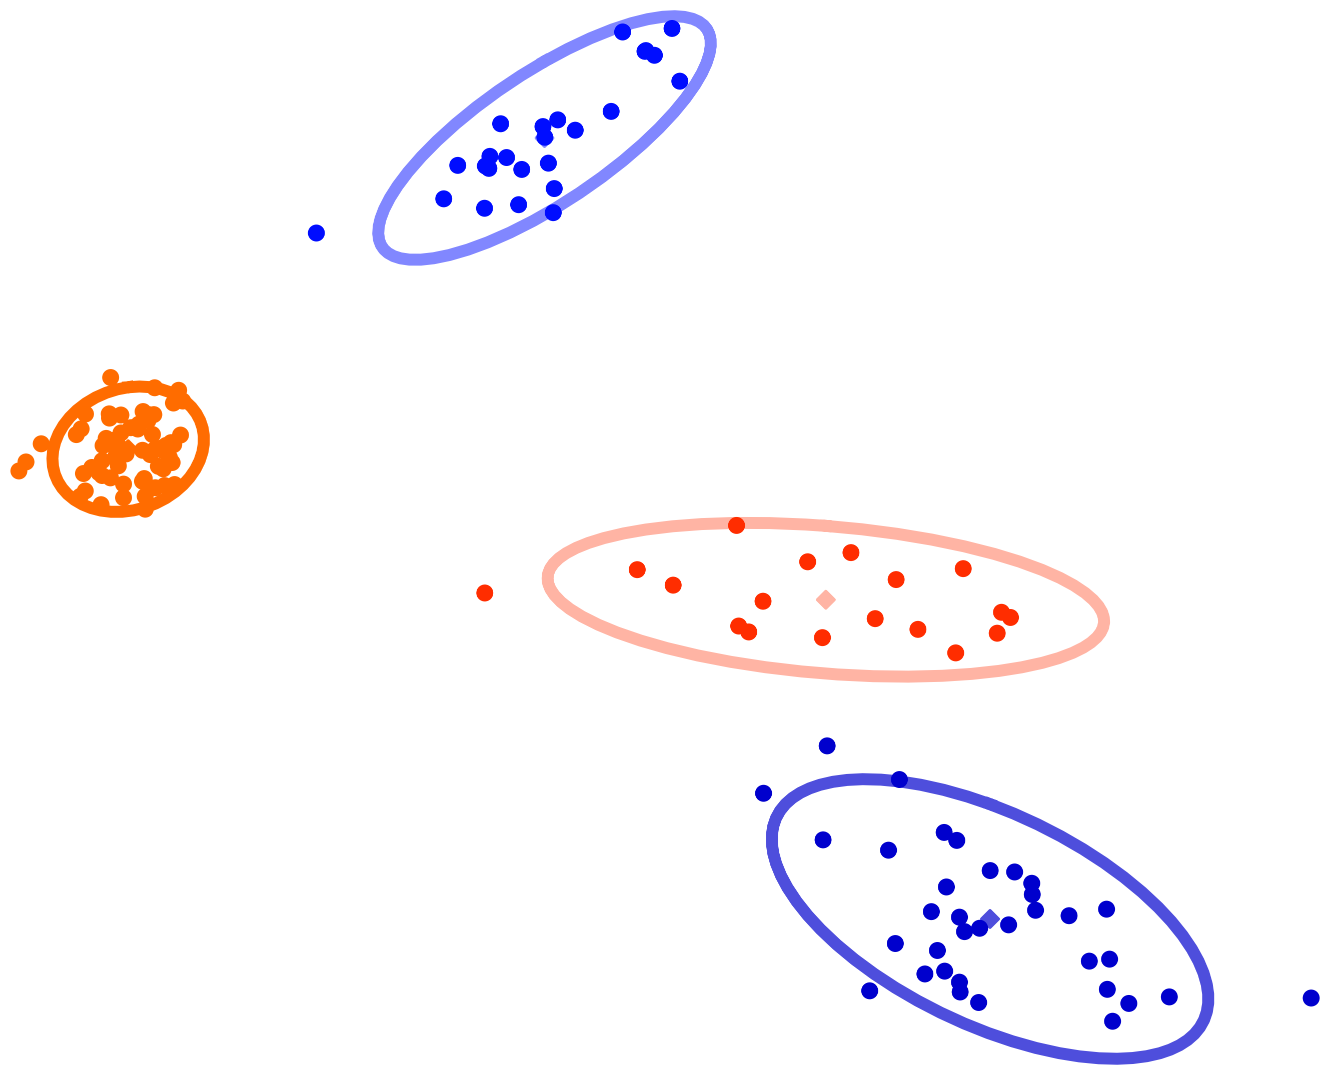

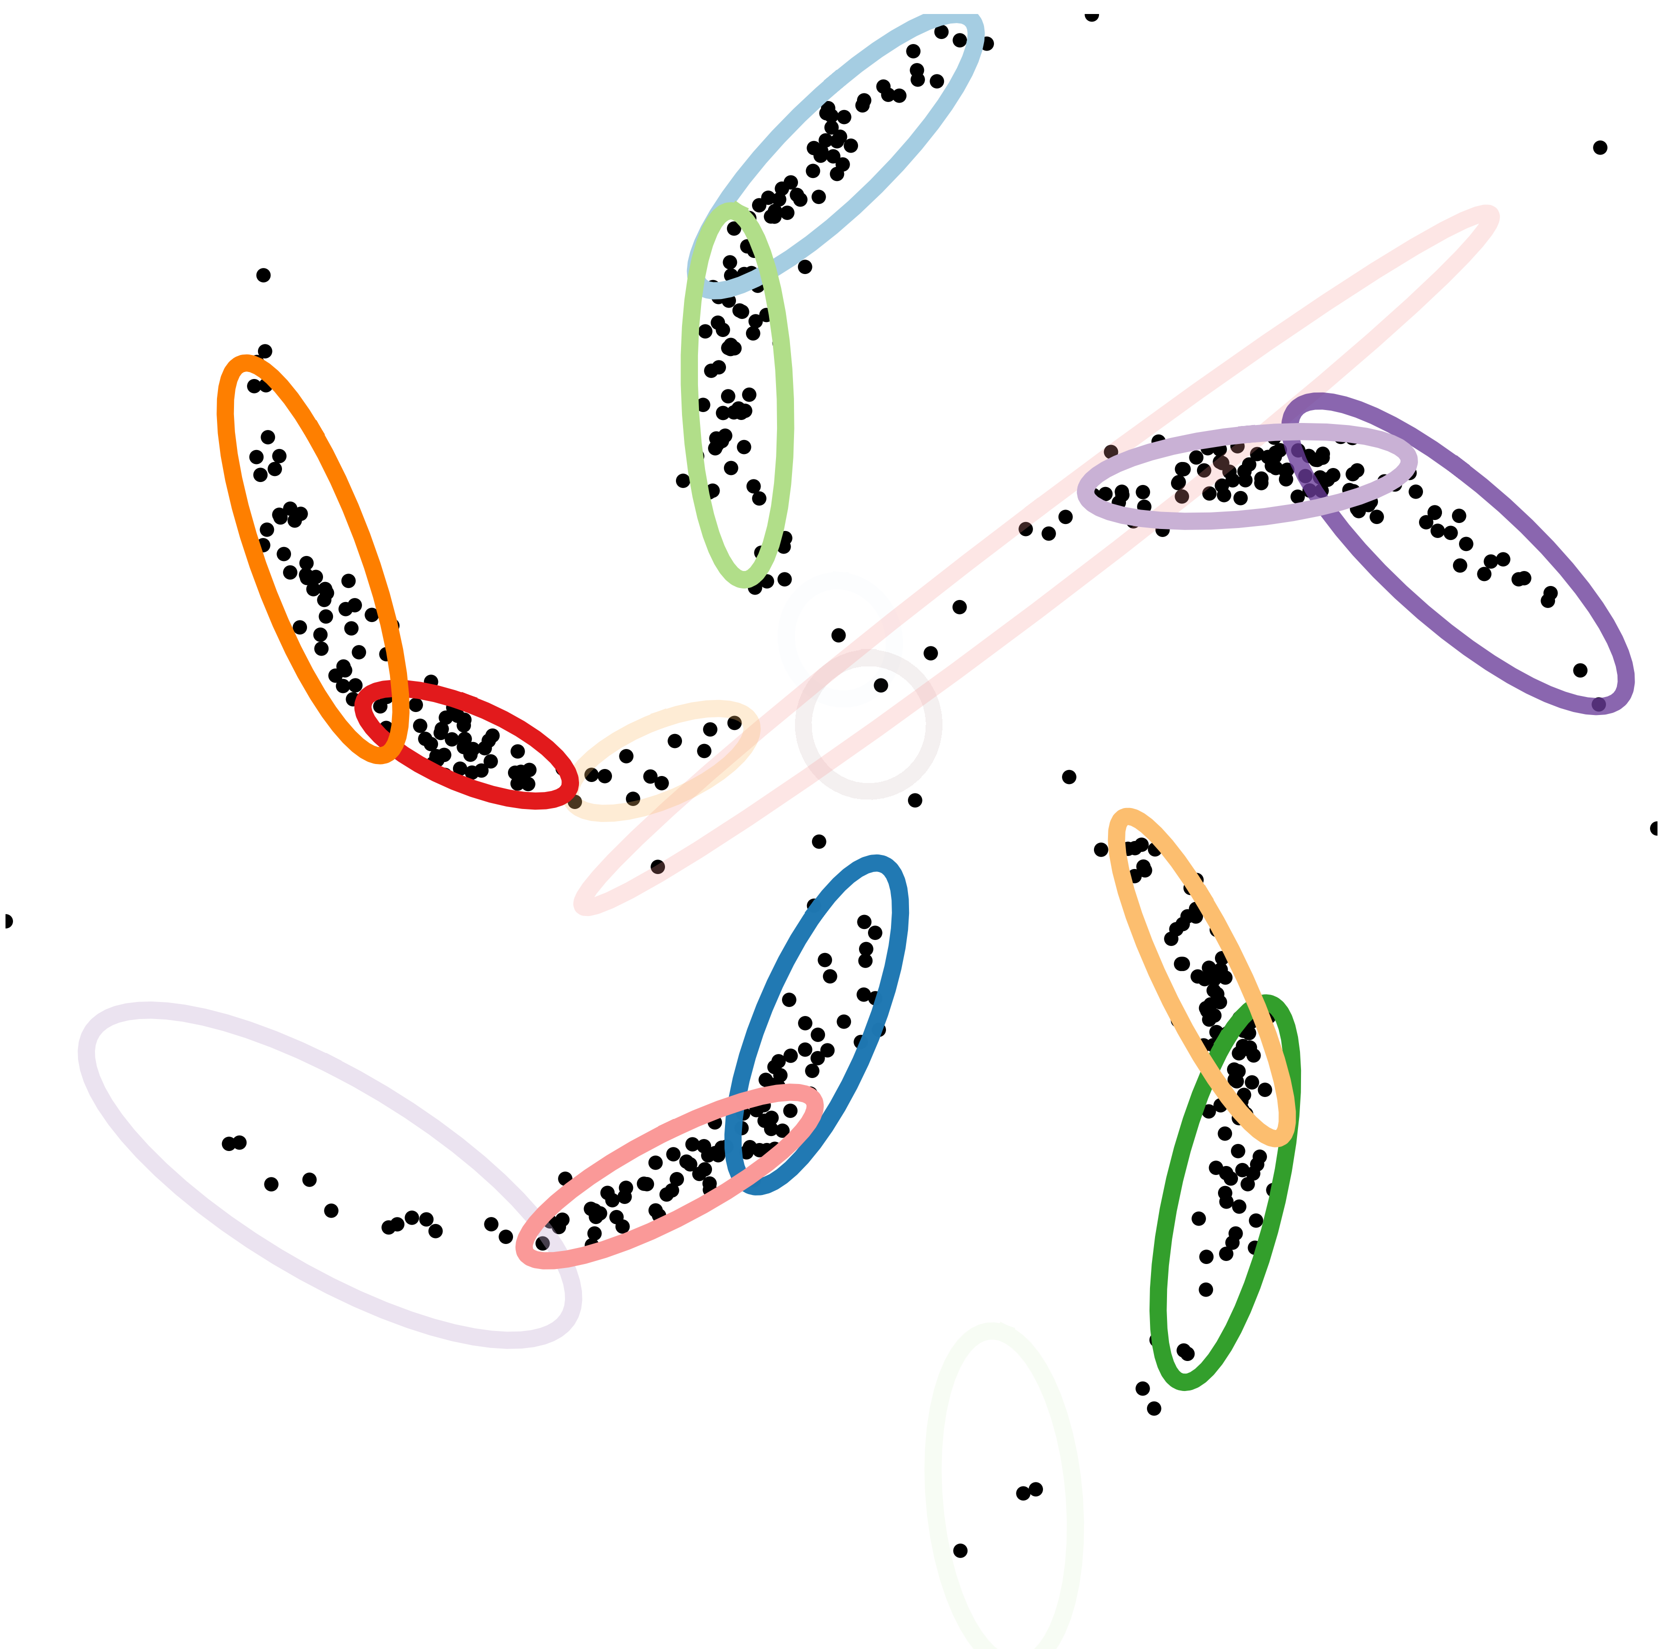

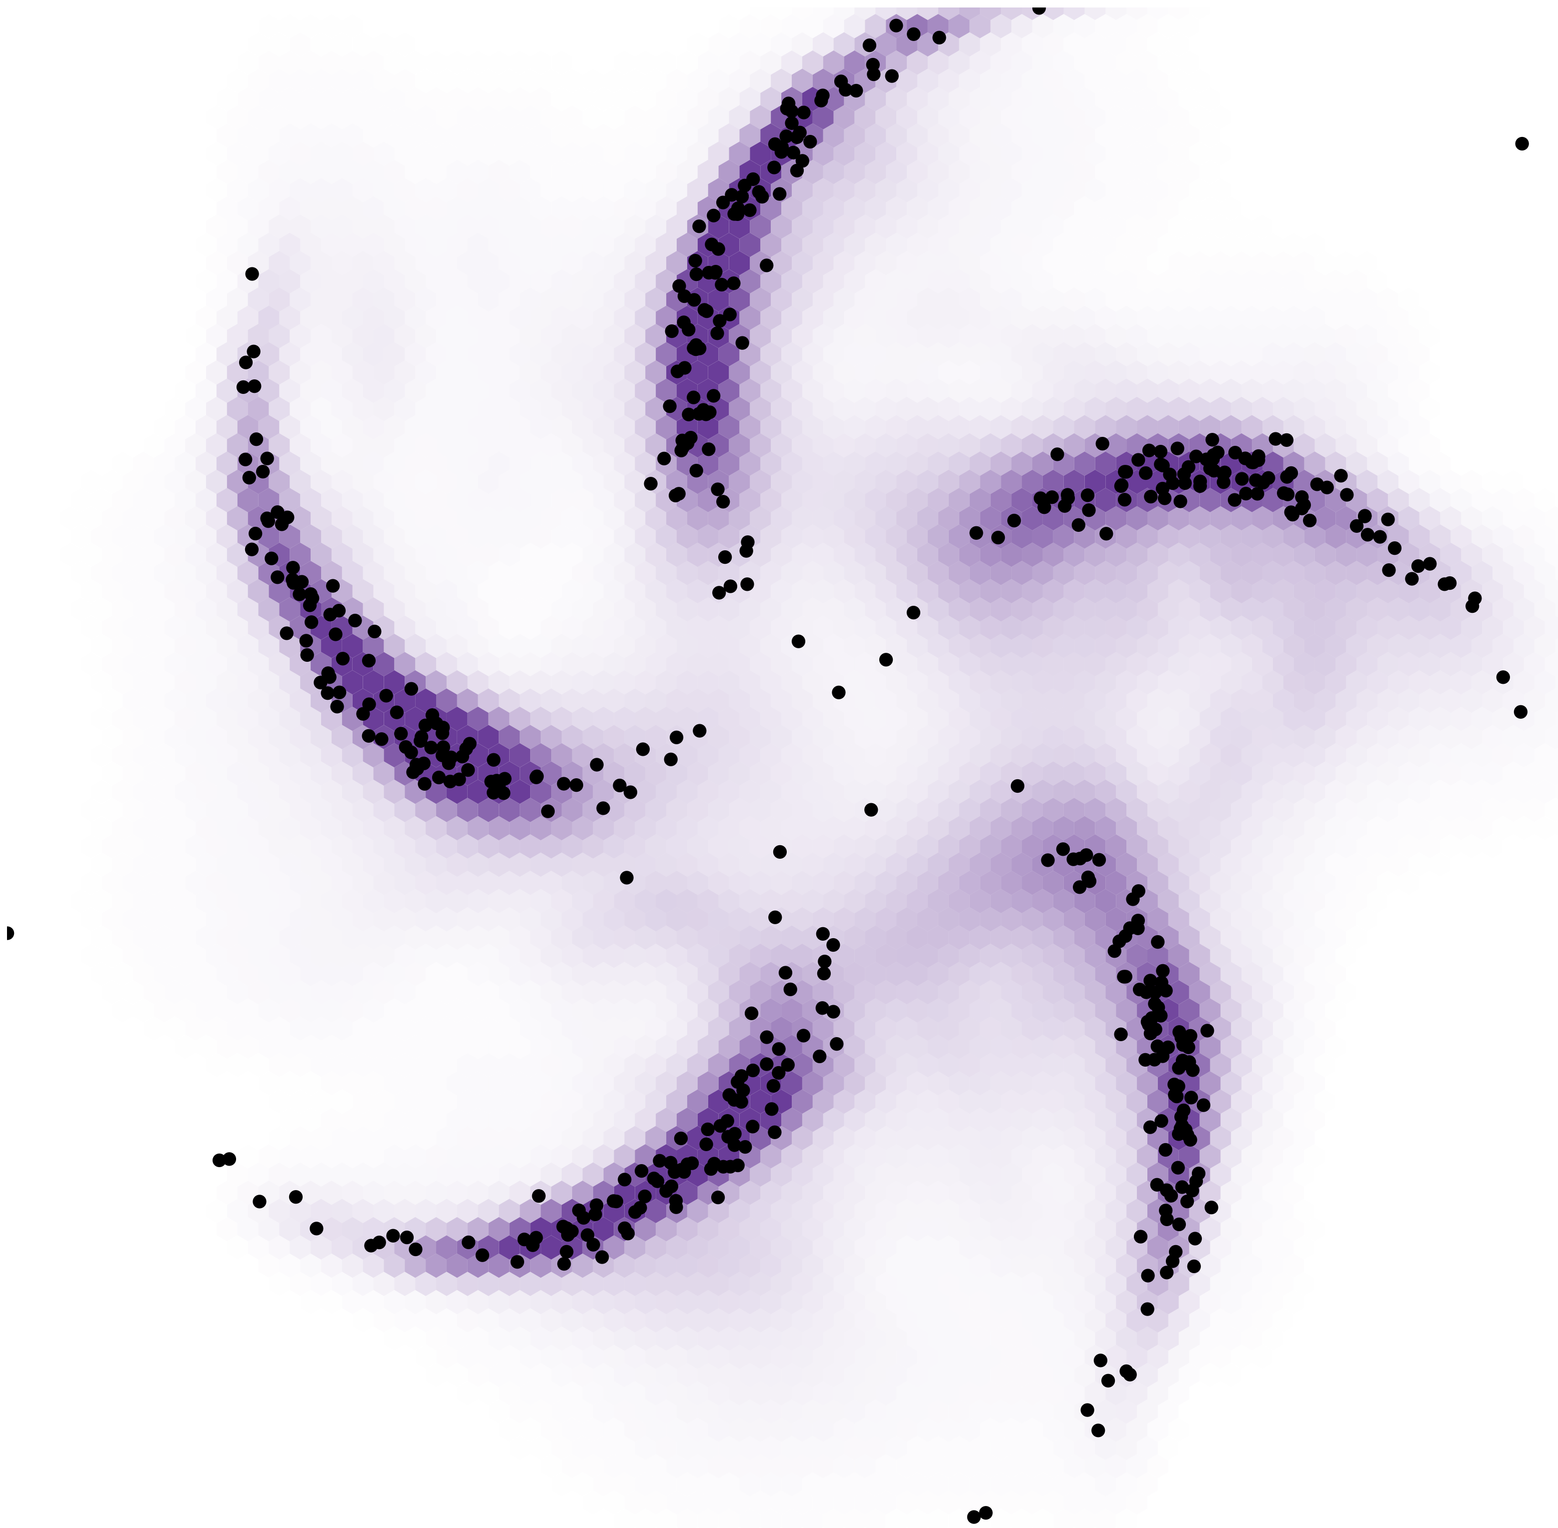

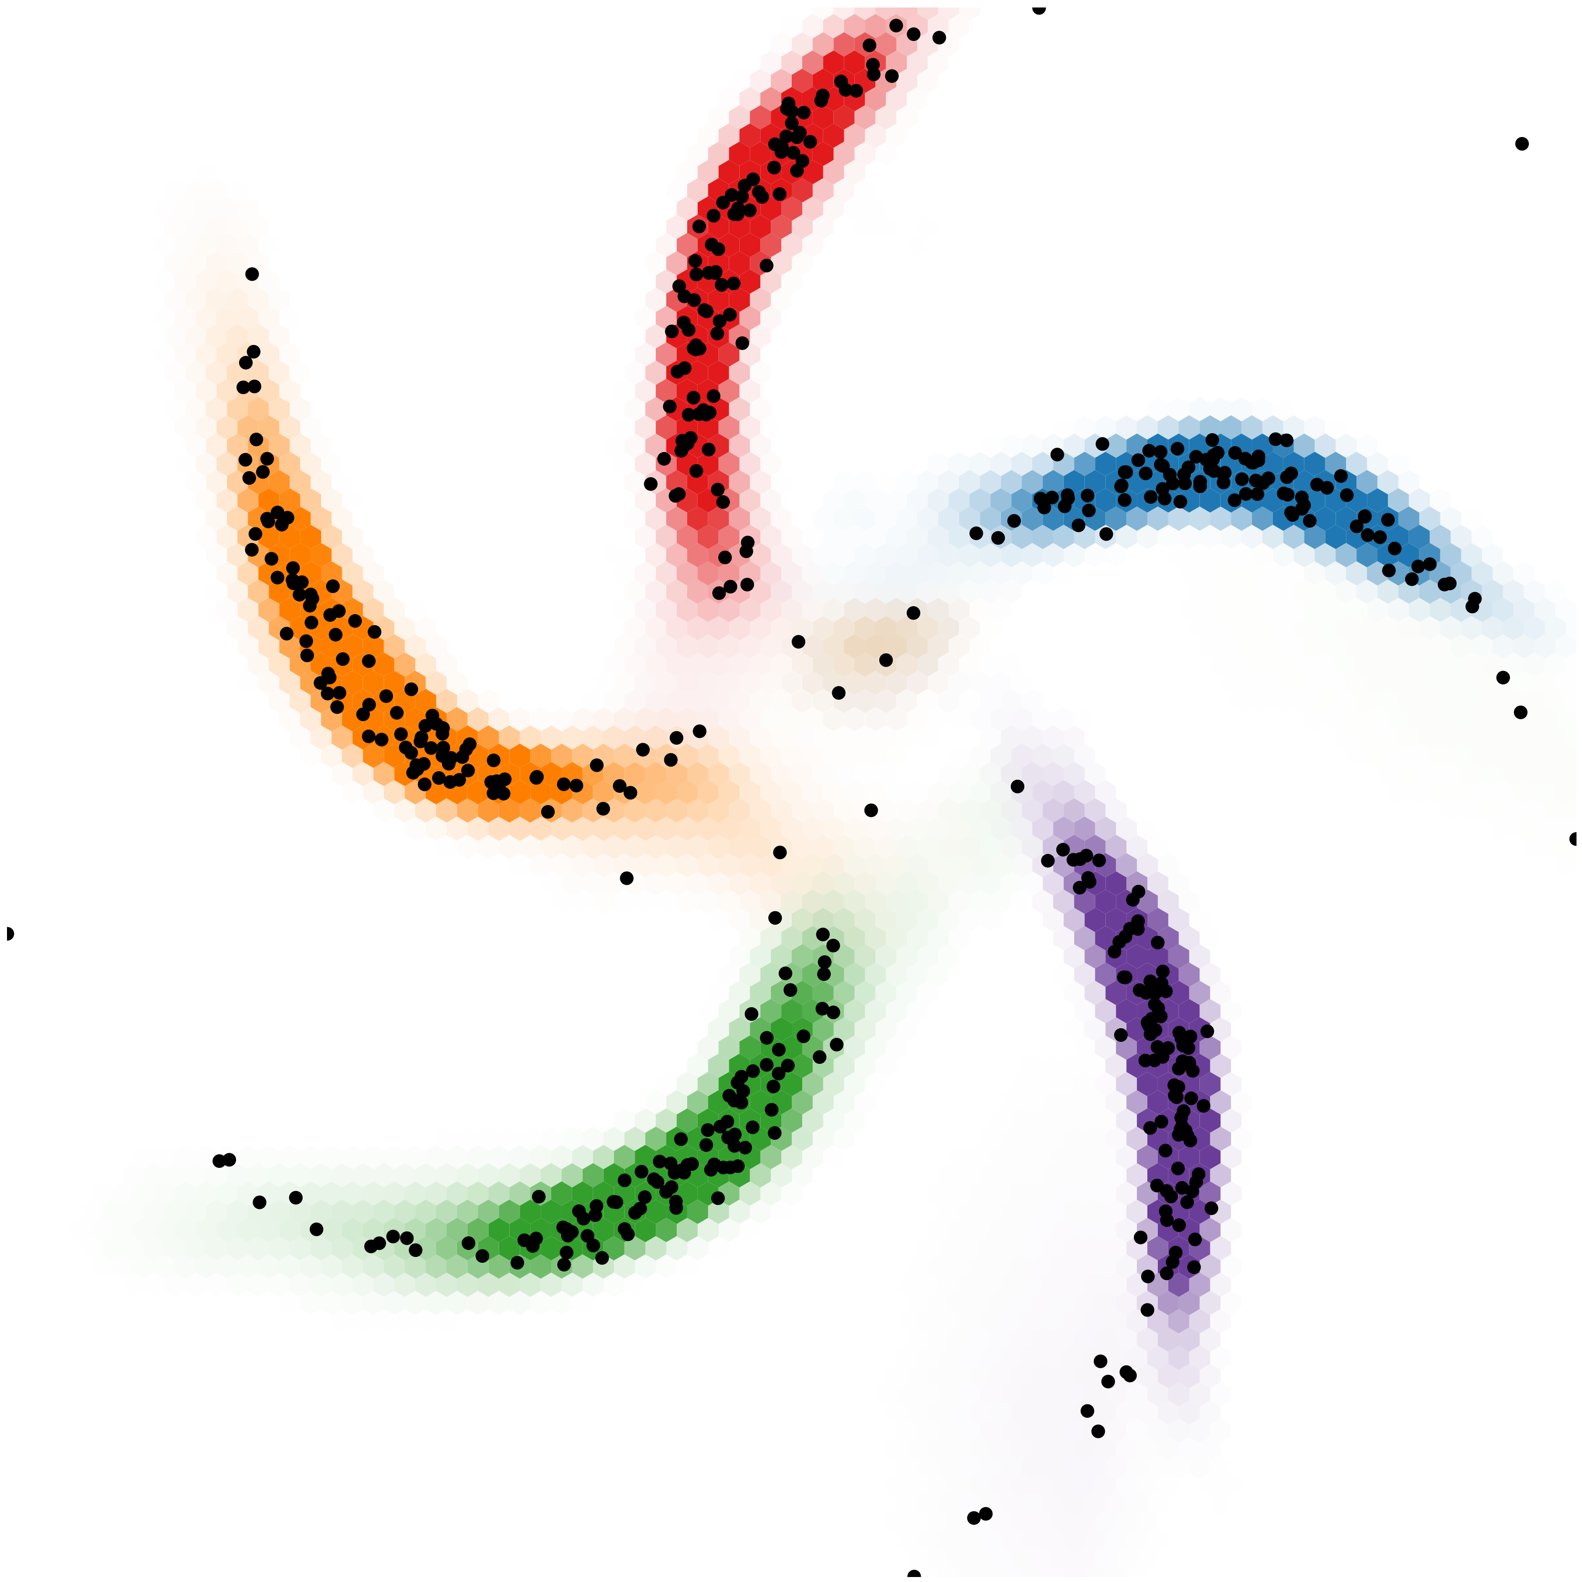

Modeling idea: graphical models on latent variables,
neural network models for observations

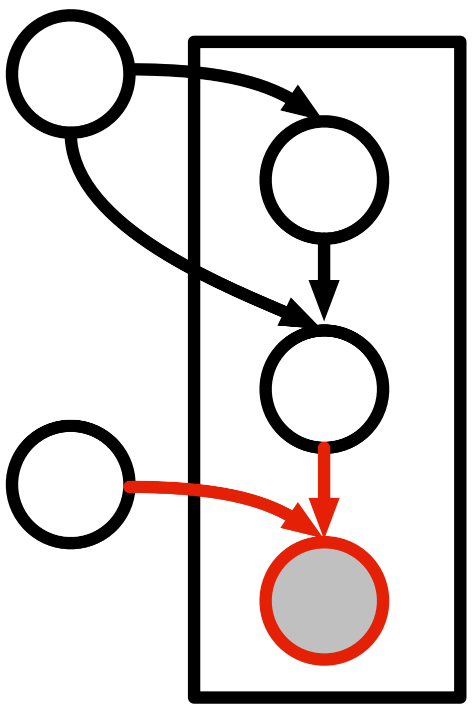

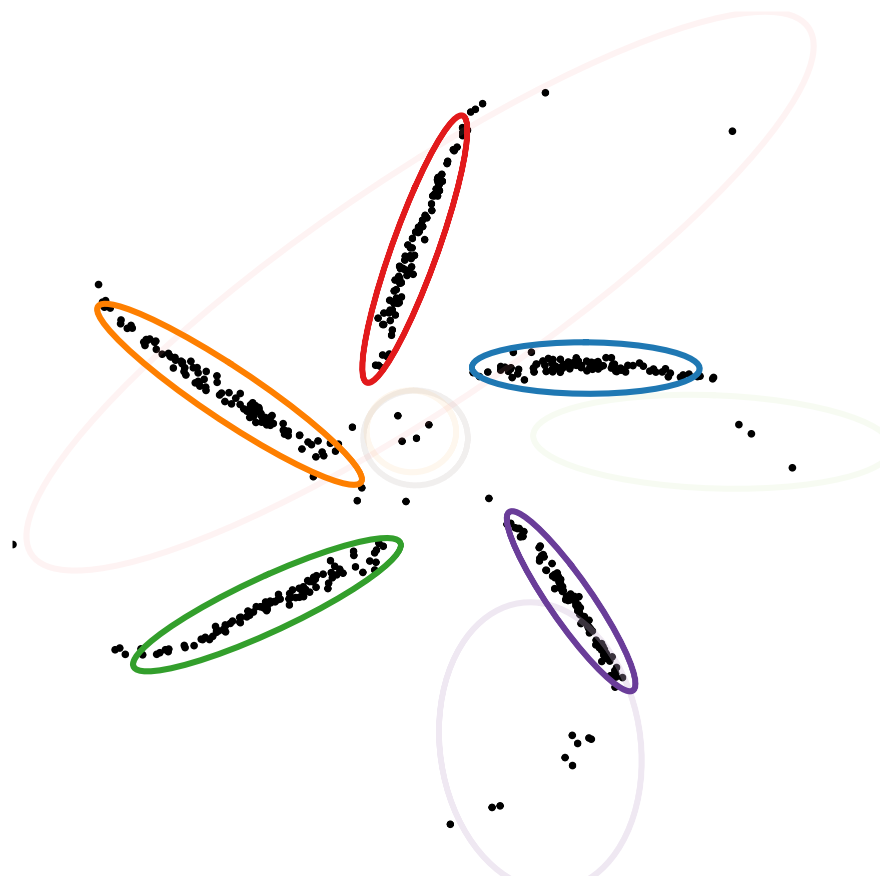

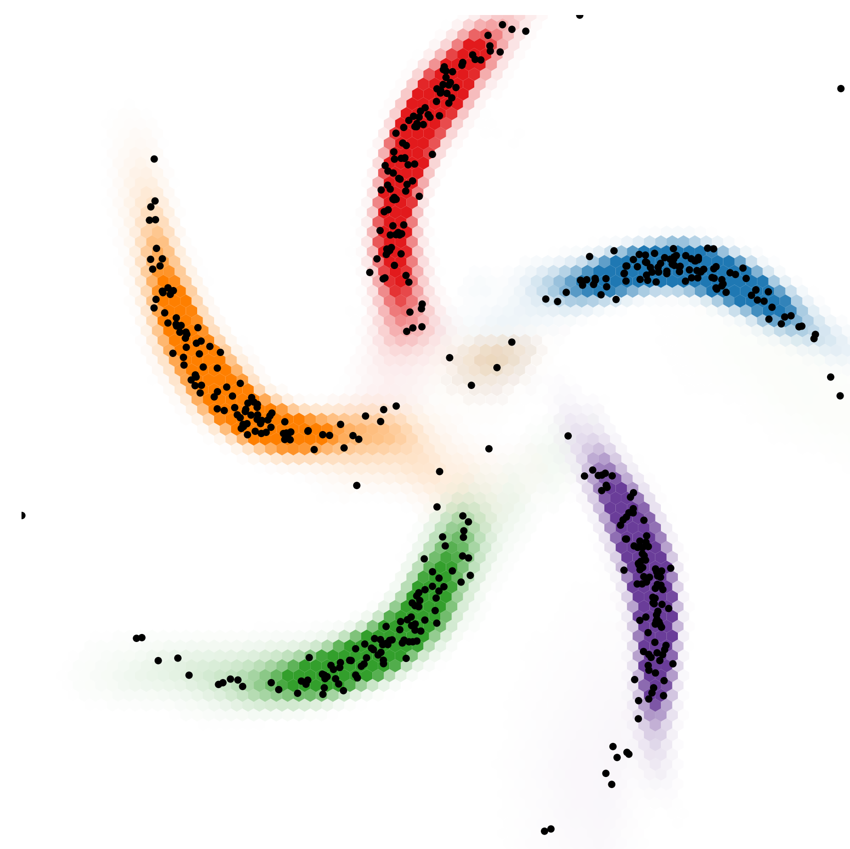

Composing graphical models with neural networks for structured representations and fast inference. Johnson, Duvenaud, Wiltschko, Datta, Adams, NIPS 2016

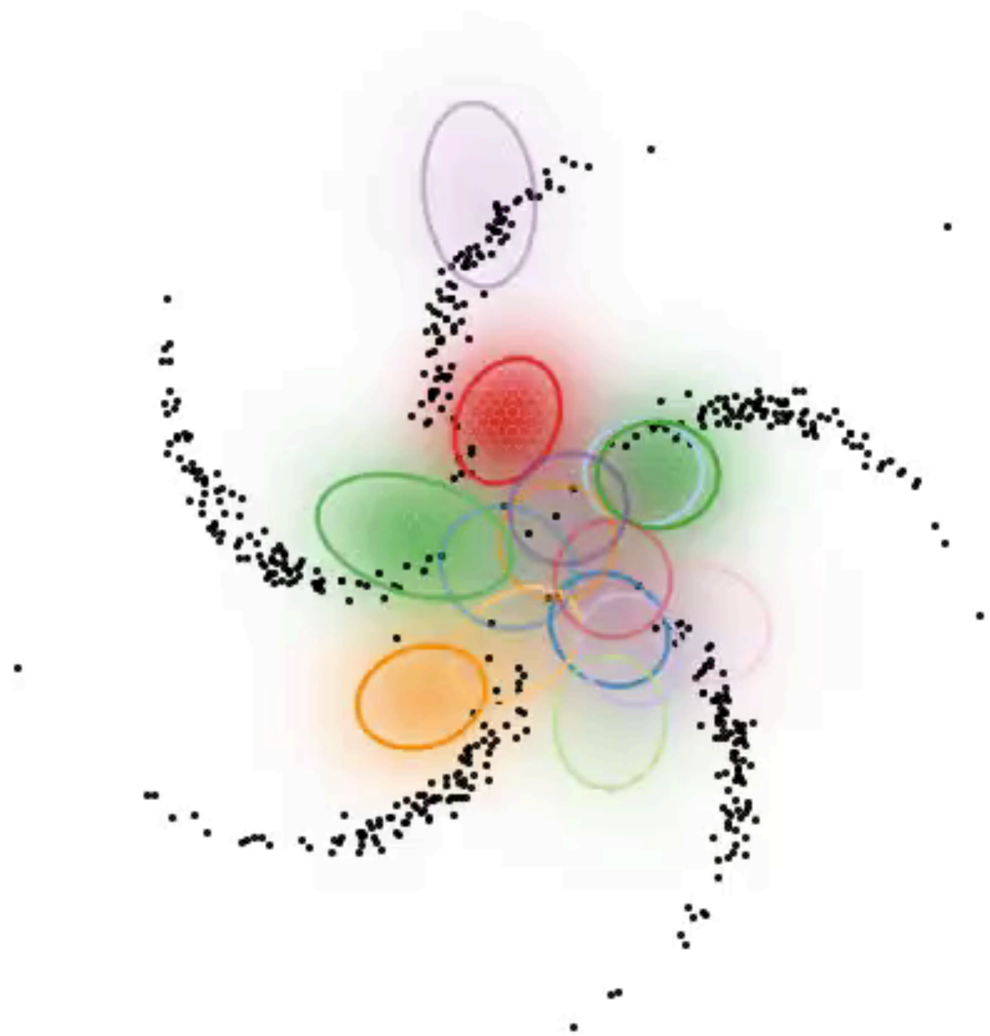

data space

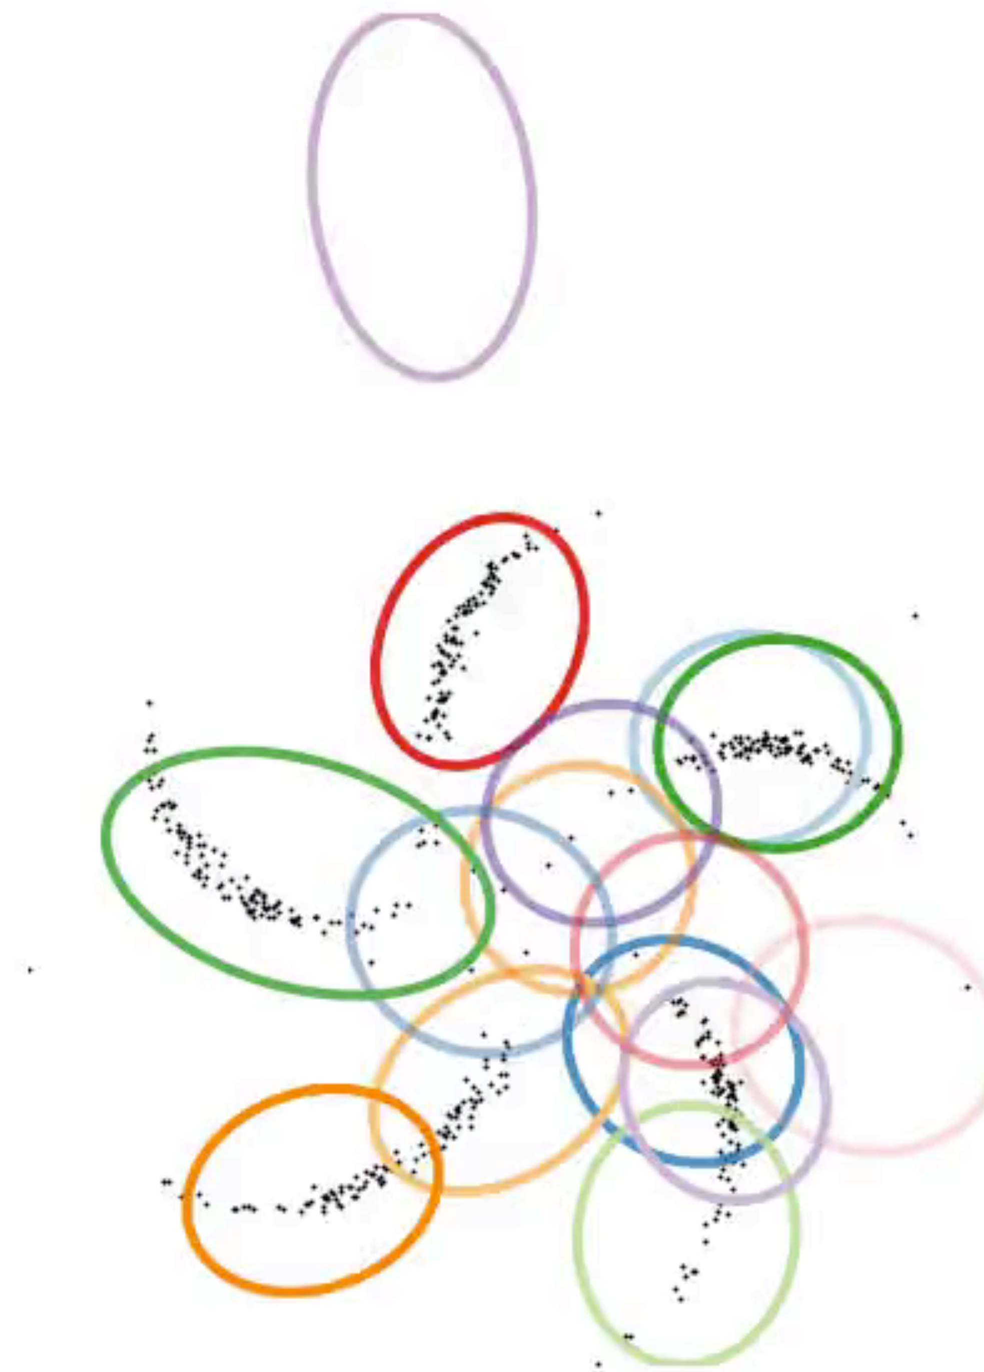

latent space

High-dimensional Bayesopt?

- Bayesian optimization doesn't really work in 50 dimensions
  - BNN instead of GP?
- No good lookahead strategies

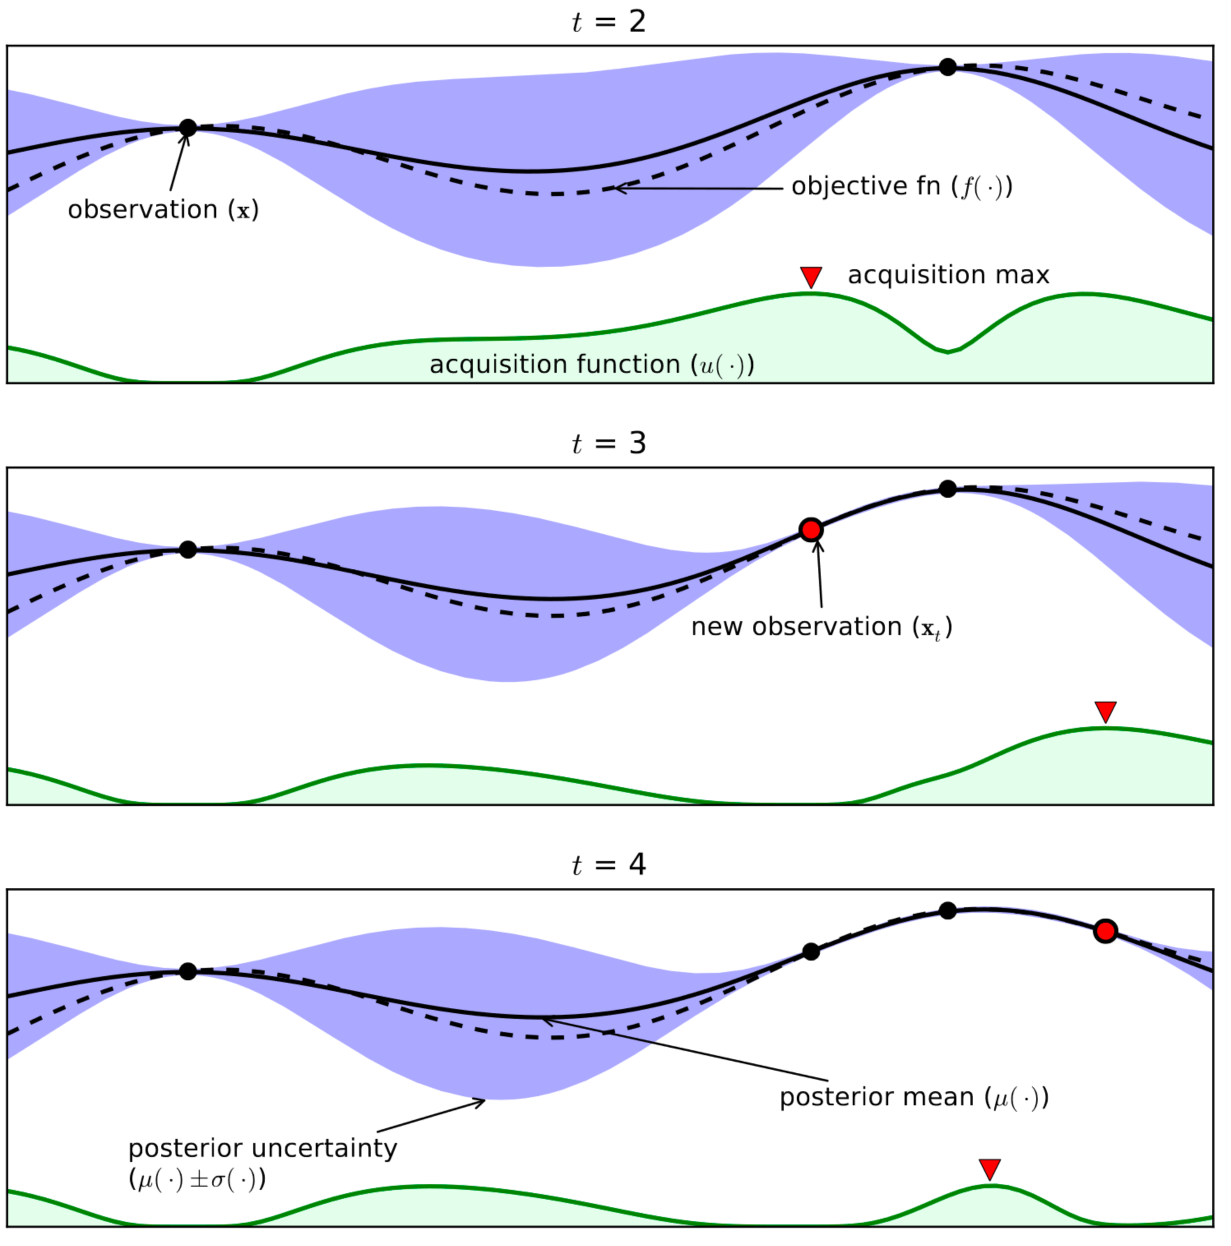

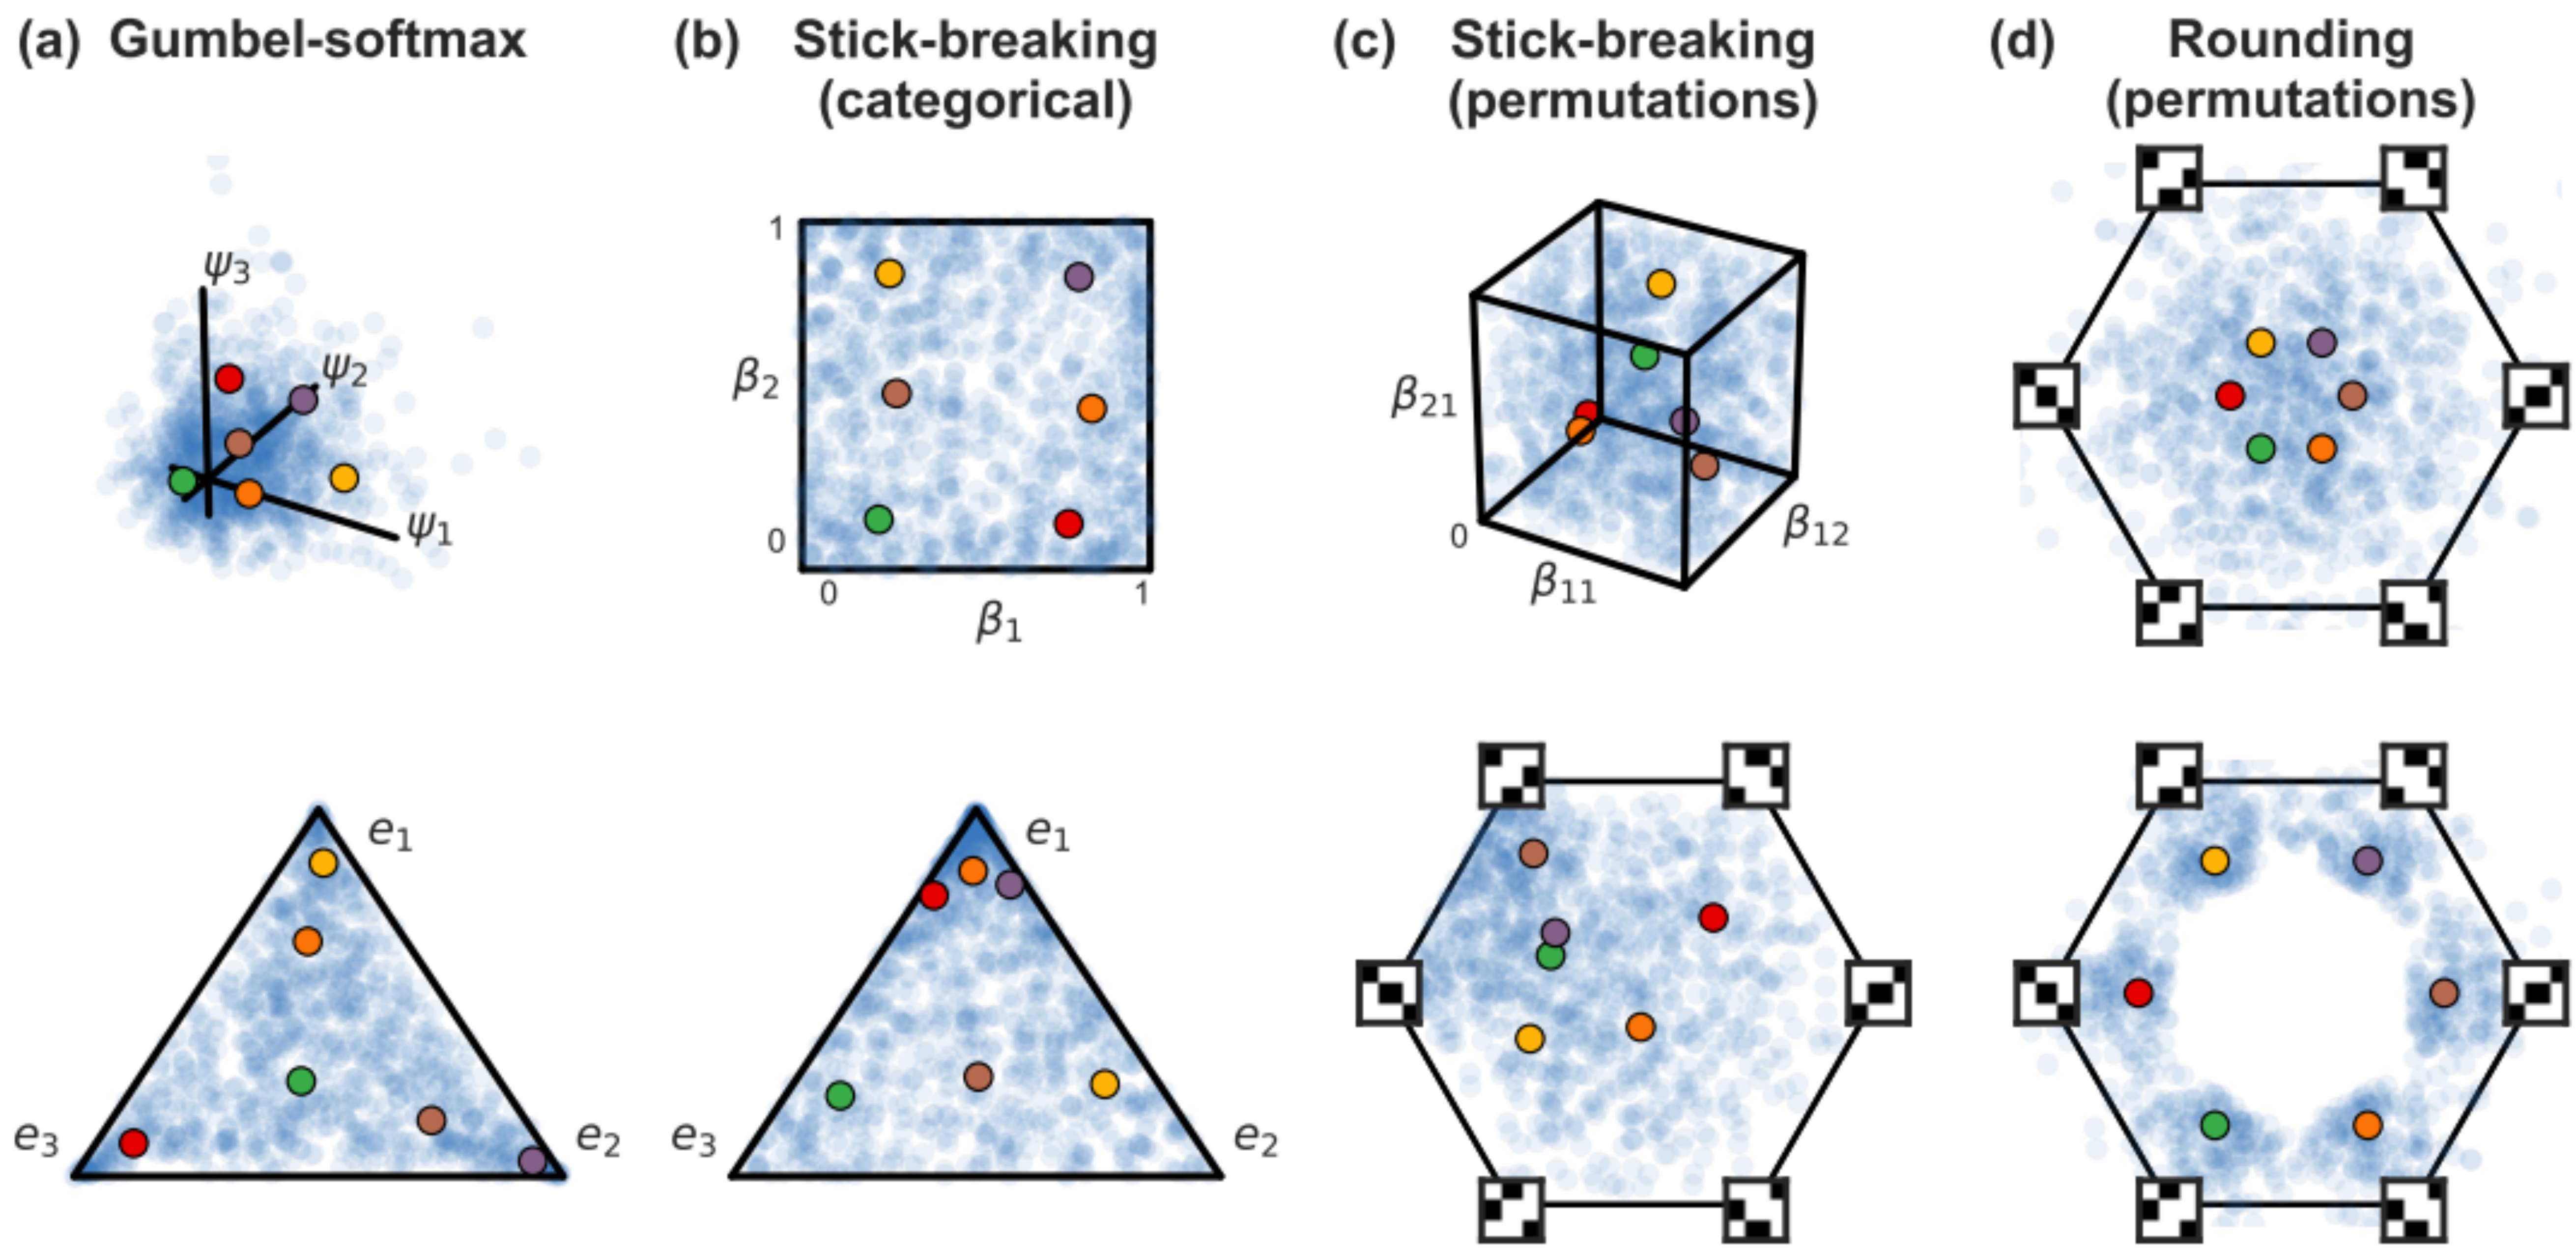

Reparameterizing the Birkhoff Polytope for Variational Permutation Inference

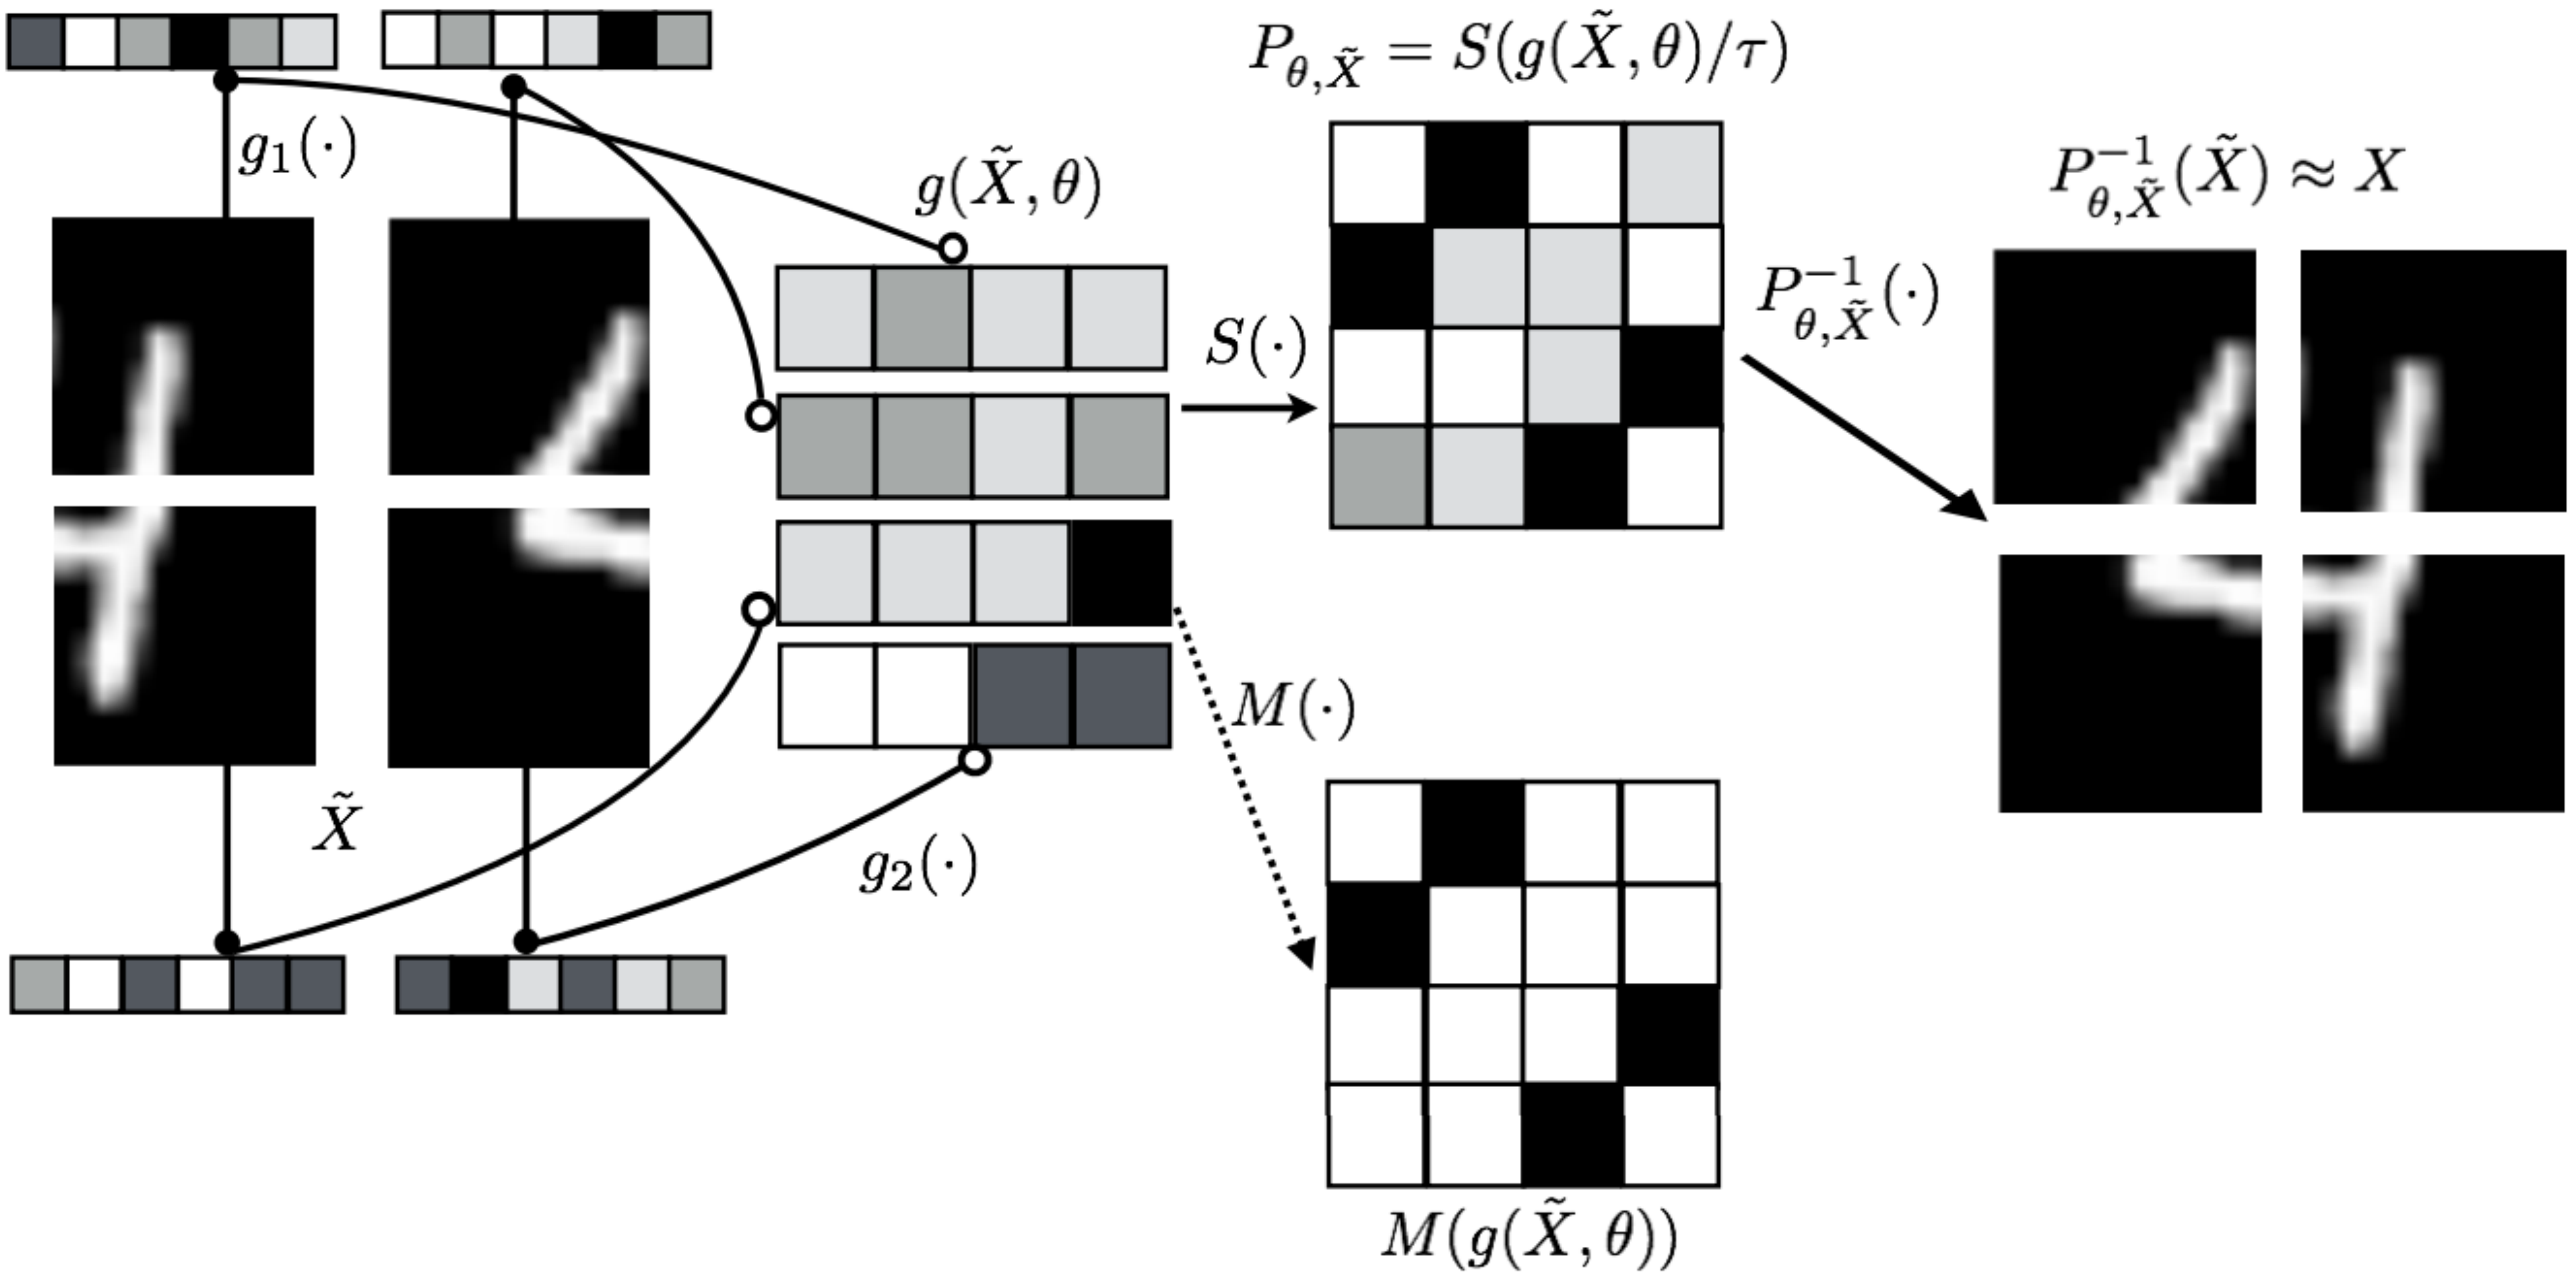

Learning Latent Permutations with Gumbel-Sinkhorn Networks

Analyzing the Surrogate

- RELAX learns to balance REINFORCE variance and reparameterization variance

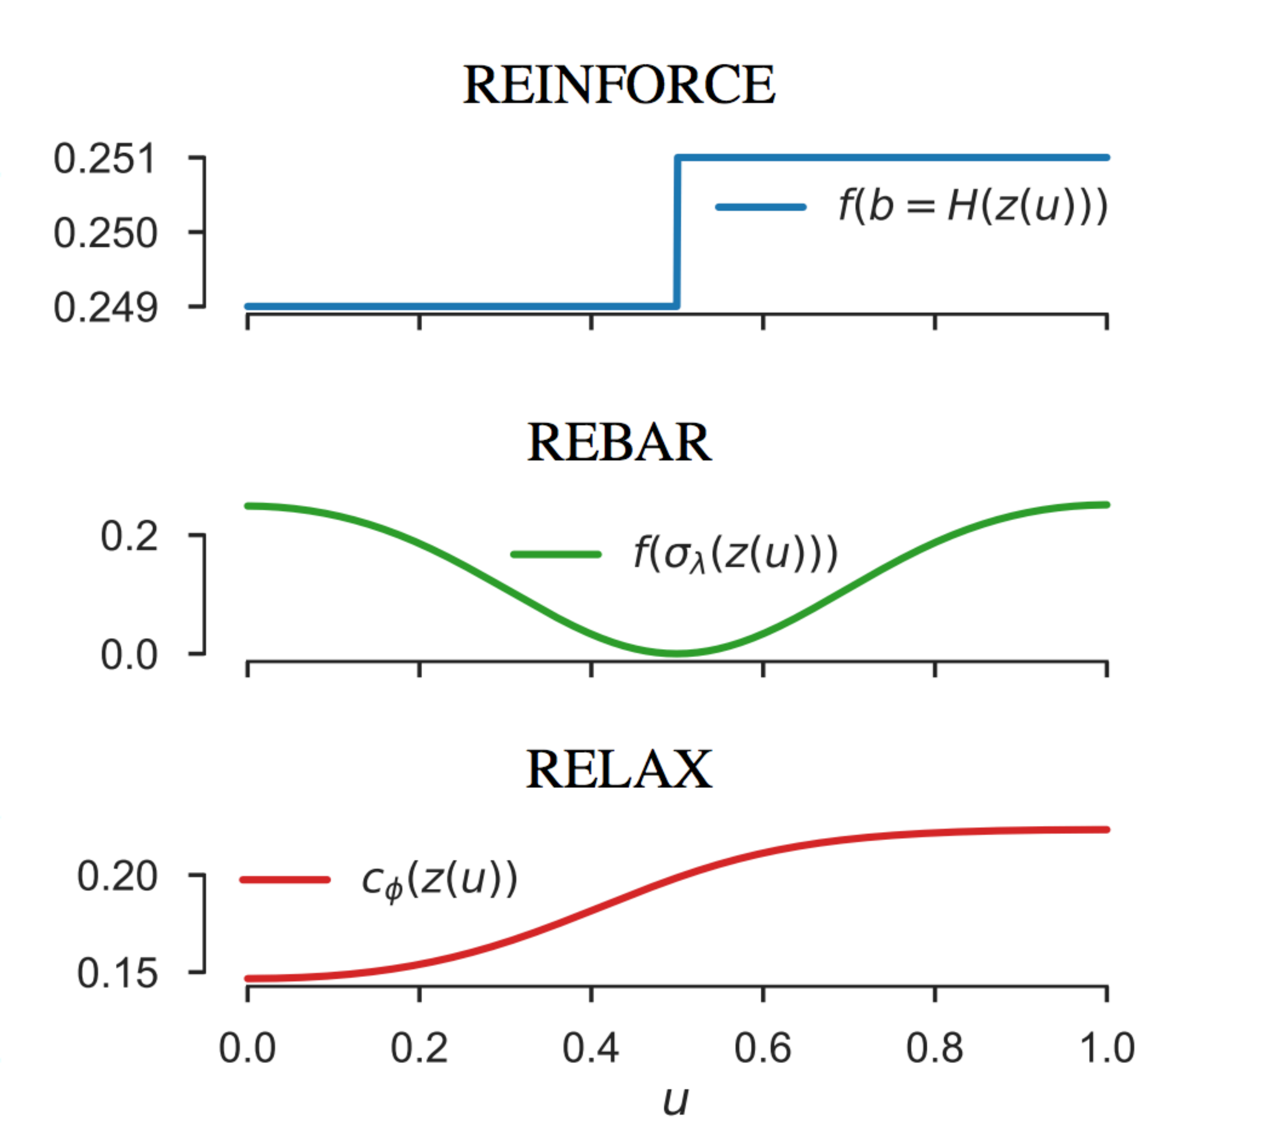

Learning to Plan

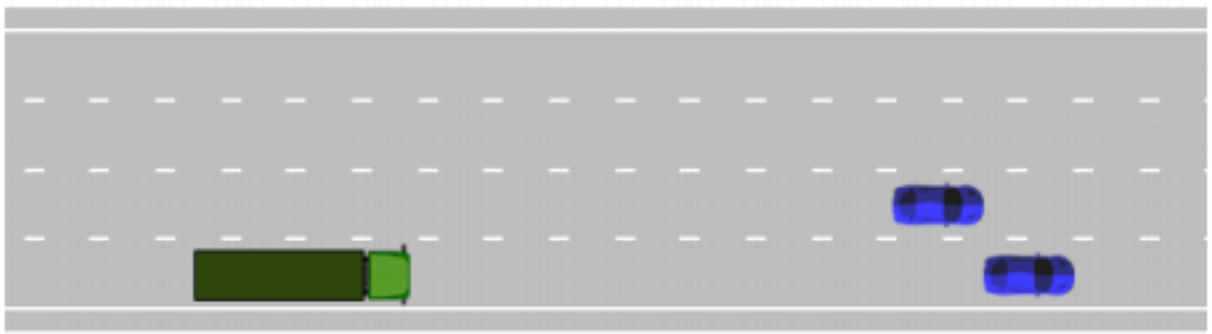

(a) $t = 0$

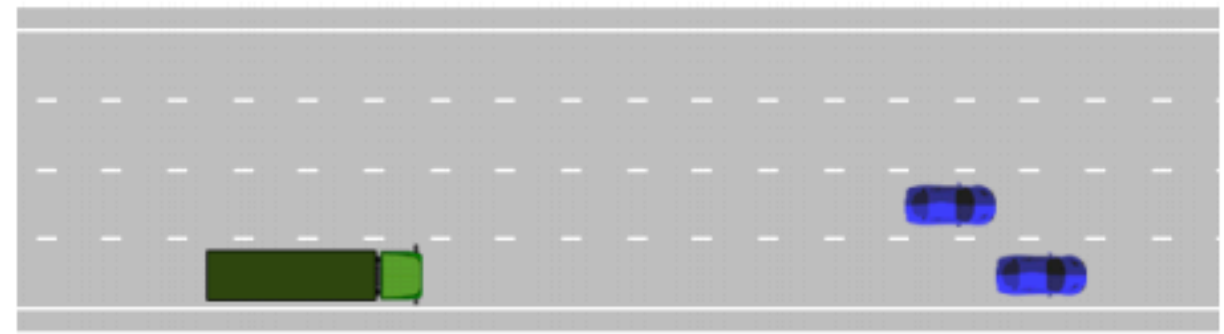

(b) $t = 15$ s, IDM/MOBIL

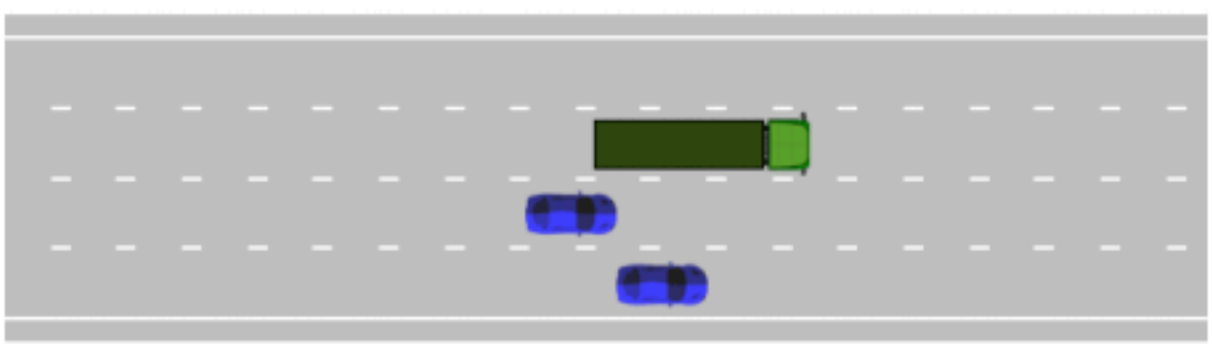

(c) $t = 15$ s, MCTS

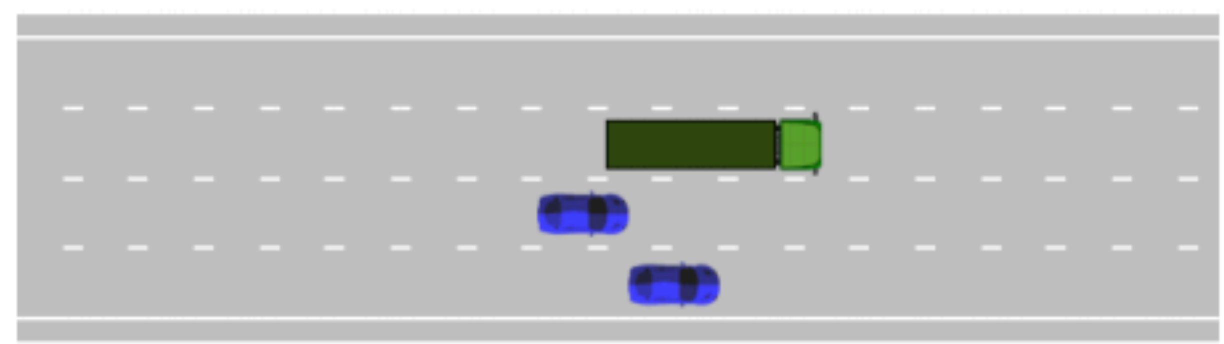

(d) $t = 15$ s, MCTS/NN

Fig. 5. Example of a situation where planning is necessary. The initial state is shown in (a) and the state after 15 s is shown for the three agents in (b), (c), and (d). The green truck is the ego vehicle.

Hoel et alia, 2019

Project ideas: Easy

- Systematically compare gradient estimators for discrete expectations. E.g. investigate scaling properties of Concrete, REBAR, RELAX with dimension of latent space
- Implement REBAR or RELAX in JAX (allows cheap per-example gradients)
- Apply existing gradient estimators to an existing problem:
  - Training GANs on text, learning to communicate
  - Search for origami instructions
- Literature review of e.g. gradient estimators, SAT solver optimizers, proof search methods

Project ideas: Easy

- Study of heuristics for genetic algorithms, with demo of virtual fishtank for “Neural Graph Evolution”.

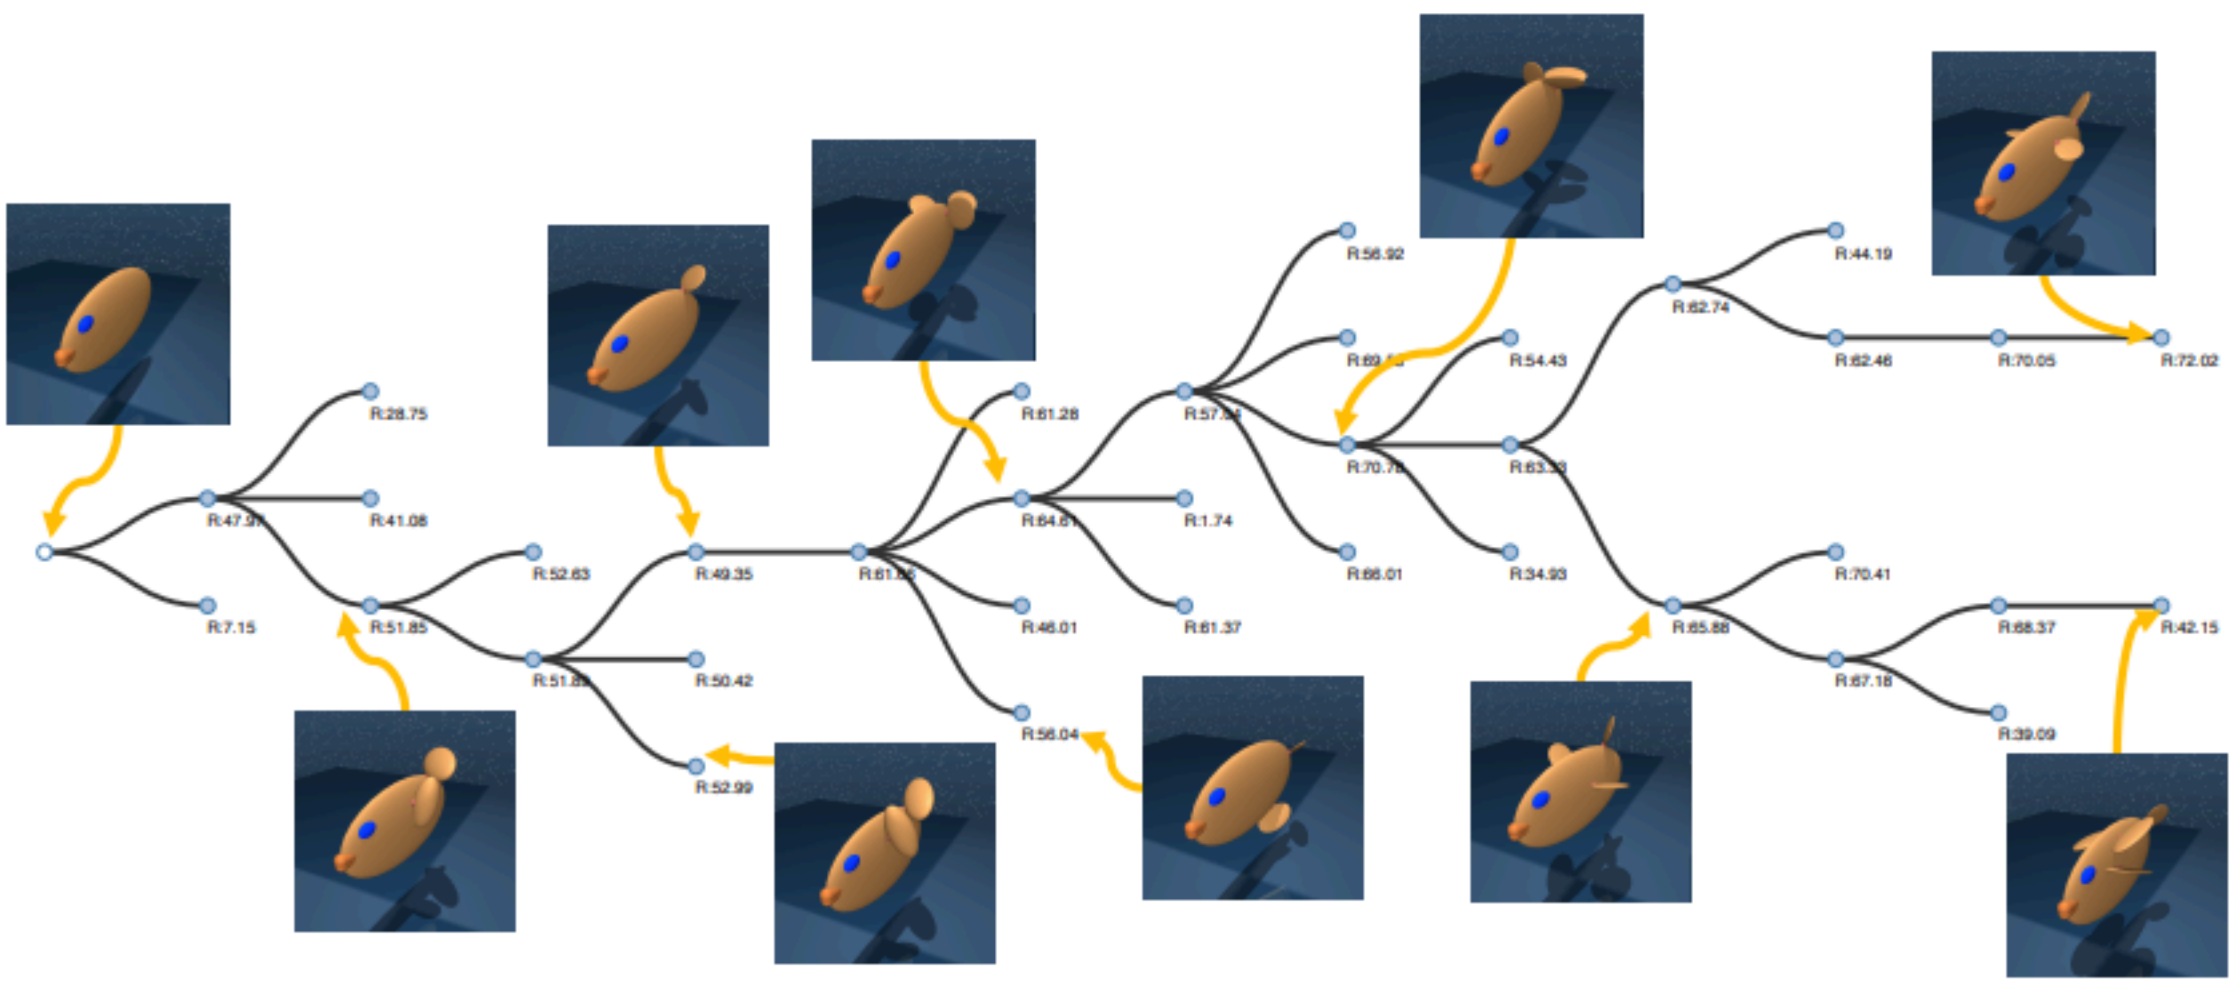

Project ideas: Medium

- Apply implicit differentiation to training GANs
  (related work ongoing by Guodong Zhang, Jimmy Ba, Roger Grosse)
- Come up with tractable approximations to K-step lookahead in active learning / search in some domain
- Learn a surrogate cost function for an existing search algorithm during the search
- Come up with a new relaxation or sampler (like Concrete, REBAR) for a new type of discrete object, e.g. permutation matrices, DAGs, hierarchies of graphs
- Regularize Deep Equilibrium Models to be easy to solve (recommended)

Project ideas: Hard

- Derive generalizations of “intrinsic motivation” and “curiosity” as approximate solutions of an MDP with distribution over rewards but known dynamics. Jeff Negrea made some progress.

- Attempt to learn tractable approximations of MDP with unknown dynamics (“learn to practice”)

- VAE for phylogenetic trees

Project ideas: Holy Grail

- Tractable approximations for solving POMDPs with unknown dynamics and rewards (i.e. simultaneous planning and learning)

- Program/proof search algorithms learns from previous and partial solutions

- General strategies for constructing low-variance gradient estimators through structured discrete variables

- Theoretical characterization of discrete optimization problems

Related (okay) Project Topics

- Continuous nested optimization: Meta-learning, recognition networks, Stackleberg games (GAN optimization), implicit differentiation

- Classic planning algorithms, active learning

Projects not in scope

- Plain supervised/unsupervised learning with continuous everything
- New continuous optimization algorithms
- Tweaking network architectures
- Applying deep learning / RL to some domain

Questions

Class Presentations

- Goal: High-quality, accessible tutorials.
- 110 students / 8 weeks = 13 students per week
- 13 students / 7 presentations per week = 2 students per presentation. Expecting good materials and clear exposition
- 2 week planning cycle:
  - Friday 2 weeks before: meet after class to divide up material
  - 7 - 10 days later: meet TA for practice presentation (required)
  - Present that Friday under strict time constraints

Draft Presentation Rubric

1. Say the first sentence of your presentation without any filler words: [5%]
2. Provide the necessary background to understand the main contribution of the paper: 20%
3. Related work: 15%
4. Explain the main ideas of the paper clearly: 20%
5. Explain the scope and limitations of the approach, or open questions 10%
6. Show a visual representation of one of the ideas from the paper: 10%
7. Original content: 10%
8. Finish under time: 5%
9. Get feedback from TAs ahead of time: 5%

Class Presentations

- Need volunteers for presenting Sept 27th on MCTS. Meet right after class, then on Monday/Tues
  - Extra support
  - Avoids overlap with assignment / project proposal / presentation
- Other weeks will be based on a sign-up survey next week
  - available to waitlisted students in case slots open up

Office Hours

- My office hours - 1h/week
- Regular TA office hours - 1h/week
- Project proposal TA office hours - 3h/week for two weeks
- Project TA office hours - 3h/week for last two weeks

Shengyan Sun

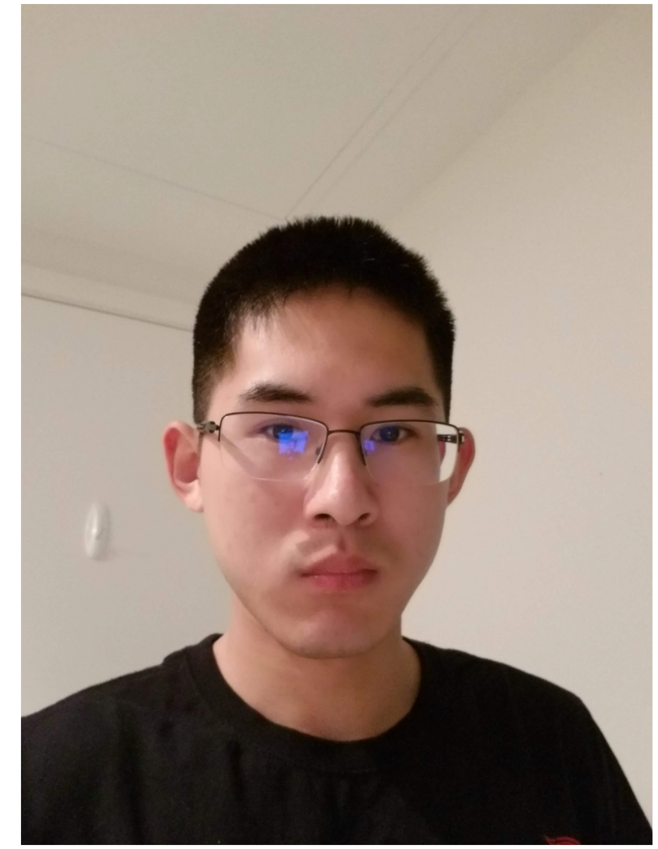

Research Interests:

- o Bayesian modelling, from both empirical and theoretical sides.
- o Reasoning with propositional and higher-order logic.
- o SAT solvers and theorem proving

1. J. Yang*, S. **Sun***, D. Roy. Fast-rate PAC-Bayes Generalization Bounds via Shifted Rademacher Processes. NeurIPS 2019.
2. **S. Sun***, G. Zhang*, J. Shi*, R. Grosse. Functional variational Bayesian neural networks. ICLR 2019.
3. **S. Sun**, G. Zhang, C. Wang, W. Zeng, J. Li, and R. Grosse. Differentiable compositional kernel learning for Gaussian processes. ICML 2018.
4. J. Shi, **S. Sun**, J. Zhu. A Spectral Approach to Gradient Estimation for Implicit Distributions. ICML 2018.
5. G. Zhang*, **S. Sun***, D. Duvenaud, R. Grosse. (2017). “Noisy Natural Gradient as Variational Inference”. ICML 2018.

Chris Cremer

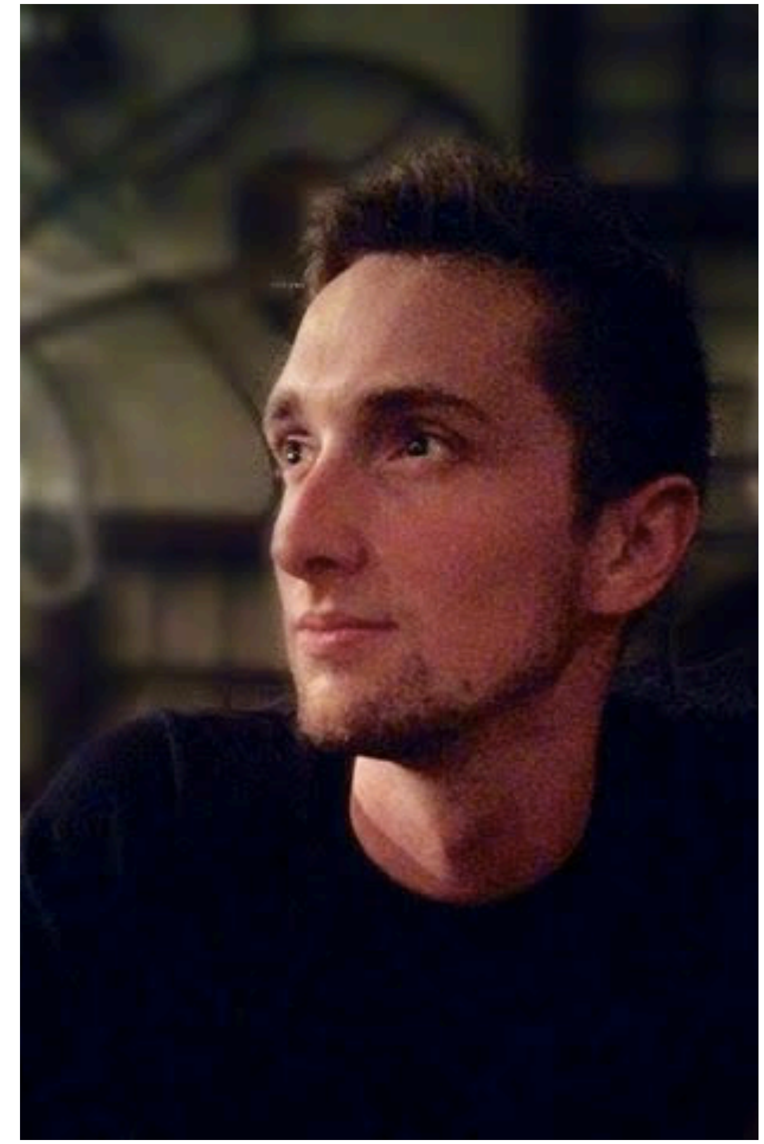

Research Interests:

- Approximate Inference
- Gradient estimation for discrete distributions
- Exploration in RL
- Model-based RL

- Cremer, C., Li, X. & Duvenaud, D. Inference Suboptimality in Variational Autoencoders. ICML 2018.

- Cremer, C., Morris, Q. & Duvenaud, D. Reinterpreting Importance-Weighted Autoencoders. ICLR Workshop 2017.

- Cremer, C. & Kushman, N. On the Importance of Learning Aggregate Posteriors in Multimodal Variational Autoencoders. AABI 2018.

Jon Lorraine

Research interests:

- meta-learning,
- learning with multiple agents,
- intersection of machine learning with game theory.

- MacKay, M., Vicol, P., Lorraine, J., Duvenaud, D., Grosse, R. Self-Tuning Networks: Bilevel
- Lorraine, J., Duvenaud, D. Stochastic Hyperparameter Optimization through Hypernets
- Lorraine, J., Hossein, S. JacNet: Learning Functions with Structured Jacobian

Quiz