Gibbs Sampling with People

Peter M. C. Harrison[1]*, Raja Marjieh[1]*, Federico Adolfi[1], Pol van Rijn[1], Manuel Anglada-Tort[1],
Ofer Tchernichovski[2], Pauline Larrouy-Maestri[1], Nori Jacoby[1]
[1]Max Planck Institute for Empirical Aesthetics, [2]Hunter College CUNY (* equal contribution)

Introduction

- A core problem in cognitive science and machine learning is to understand how humans derive semantic representations from perceptual objects [1,2,3,4].
- Markov Chain Monte Carlo with People (MCMCP) is a prominent method for studying such representations [1]. It is asymptotically powerful but can be inefficient because it relies on binary choices.
- We therefore propose Gibbs Sampling with People (GSP), a new adaptive technique based on continuous sliders.
- We reformulate MCMP and GSP in terms of utility theory and show that both methods produce sampling distributions that are proportional to the underlying utility.
- We conduct 16 experiments evaluating GSP in four domains - colors, musical chords, vocal emotions, and faces - and show that it delivers novel and interpretable results.
- The face experiment combines GSP with a state-of-the-art image synthesis network, StyleGAN, and a recent network interpretability technique, GANSpace [5], enabling GSP to efficiently explore high-dimensional perceptual spaces.

From Samplers to Humans

The central idea behind MCMCP and later GSP is to leverage some of the powerful sampling algorithms from machine learning and physics in order to characterize semantic representations by putting humans 'in the loop'.

Markov Chain Monte Carlo with People (MCMCP)

Markov Chain Monte Carlo with People [1] is based on the symmetric Metropolis-Hastings sampling algorithm. This algorithm is defined as a Markov process: let $\pi(z)$ be a probability distribution from which one would like to sample, then,

1. Sample a candidate z^* from a symmetric proposal distribution $q(z^*; z)$
2. Accept candidate according to the Barker acceptance function $A(z^*; z) = \pi(z^*)/(\pi(z^*) + \pi(z))$

In an MCMCP task, the acceptance step is replaced by a human participant. The participant is asked to pick which of two alternatives better matches a category C, resulting in MCMC sampling from the participant's internal semantic distribution $\pi(z|C)$.

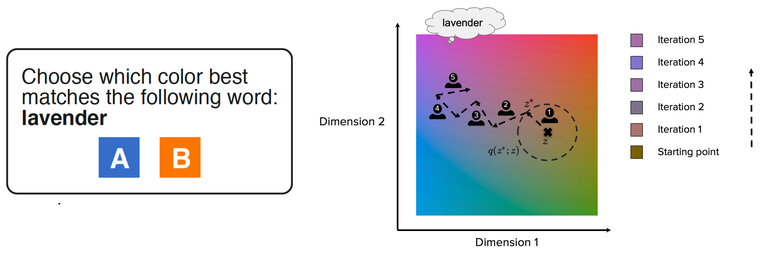

Gibbs Sampling with People (GSP)

Our novel paradigm is a human instantiation of a Gibbs sampler (see also [6] for a related discrete paradigm). To sample from a target distribution $p(\mathbf{z})$ a Gibbs sampler is defined as follows:

1. Choose a starting point $\mathbf{z}^{(1)} = (z_1^{(1)}, \ldots, z_n^{(1)})$;
2. Circularly update coordinates by sampling from $p(z_k^{(i+1)} | z_1^{(i+1)}, \ldots, z_{k-1}^{(i+1)}, z_{k+1}^{(i)}, \ldots, z_n^{(i)})$.

In a GSP task, we parametrize the space of stimuli using a set of parameters $\mathbf{z}$ of a generative model, such as a synthesizer or a neural network. At each iteration, all dimensions but one are frozen and a participant is asked to manipulate the active dimension using a slider, such that the resulting stimulus best matches a category or a criterion in question.

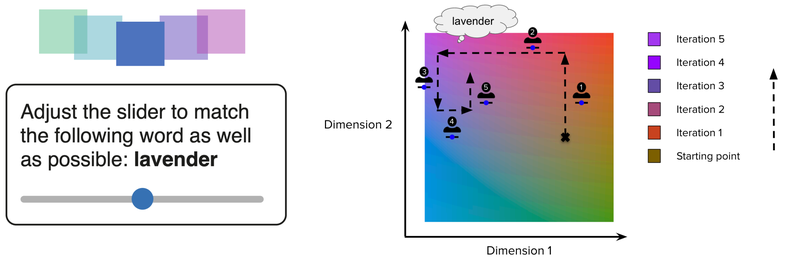

Utility Theory – Modelling Participants' Choices

- The original MCMCP formulation models the acceptance step using categorical Bayesian inference. While elegant, this does not apply to continuous semantic representations (e.g. attractiveness), and it further assumes probability matching.
- We therefore reformulate MCMCP and GSP using utility theory. According to this framework, participants extract a utility value for each stimulus $U(z) = \ell(z) + n$, with $\ell(z)$ and n being a deterministic term and stochastic noise component, respectively.
- Participants then pick the stimulus with the maximum utility. For i.i.d Gumbel noise, common to discrete choice models [7], MCMCP and GSP then sample from the exponentiated utility

$$P(\text{choose } s) = P(\forall s' \neq s : U_s > U_{s'}) = \frac{e^{\gamma \ell_s}}{\sum_{s'} e^{\gamma \ell_{s'}}}$$

- By aggregating responses from multiple participants at each iteration we can reduce the noise term and interpolate between stochastic sampling and deterministic optimization.

Aggregation

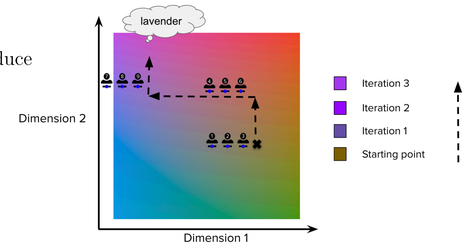

Large-Scale Online Experiments

We implemented all experiments in PsyNet, our under-development framework for implementing complex experiment paradigms such as GSP and MCMCP on MTurk. This framework builds on the Dallinger platform for experiment hosting and deployment [8].

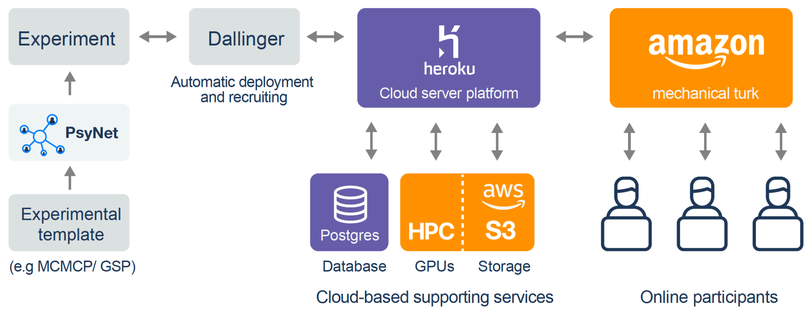

Experimental Results

Colors (3 experiments N=422, validation N=322)

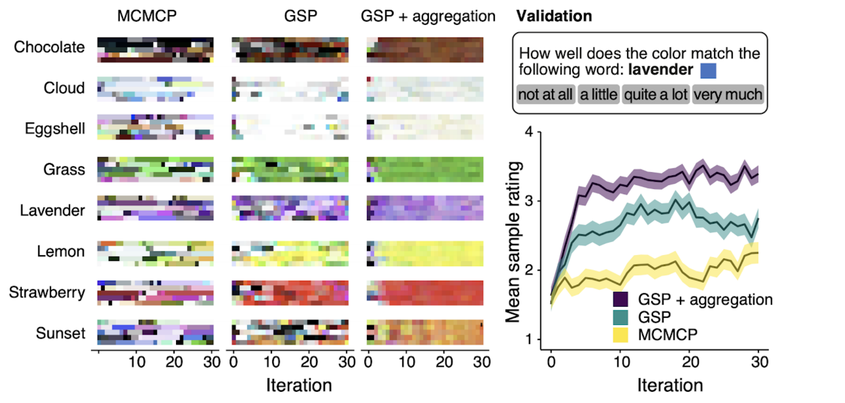

- We compare GSP and MCMCP in a toy domain - colors. We find that GSP converges on good samples faster than MCMCP, and that aggregation further improves sample quality. We validate these results using large-scale rating experiments

Chords (1 experiment N=134, validation N=168)

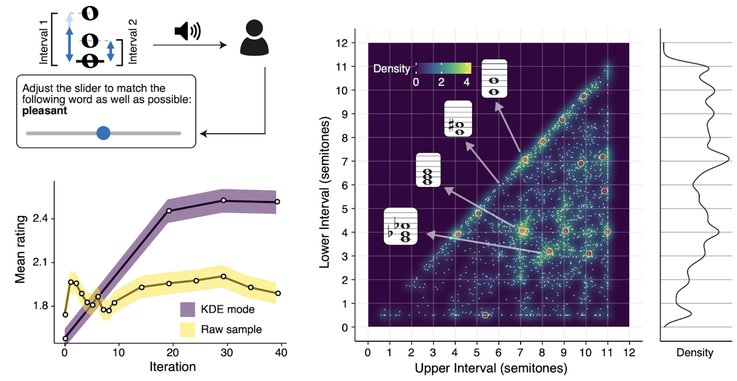

- We then apply GSP to 3-note musical chords, parametrized as pairs of continuous pitch intervals. The process reveals a rich landscape of triadic sonorities and intervallic hierarchies embedded within the Western 12-tone scale.

Prosody (2 experiments N=167, validation N=314)

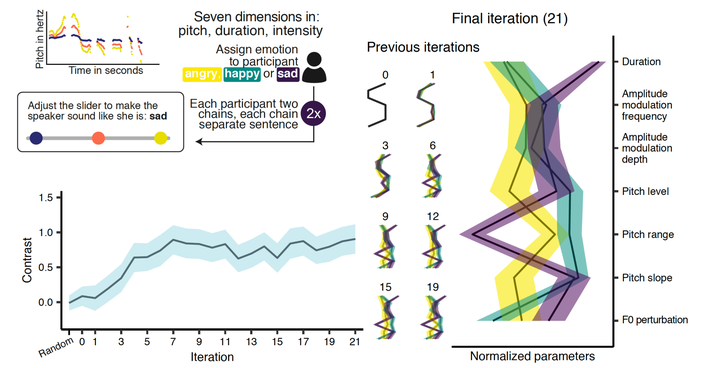

- In our next study, participants manipulated sentences along 7 dimensions to match three emotions. The process quickly identified characteristic fingerprints of each emotion that can be used to synthesise new recordings.

Faces (3 experiments N=840, validation N=440)

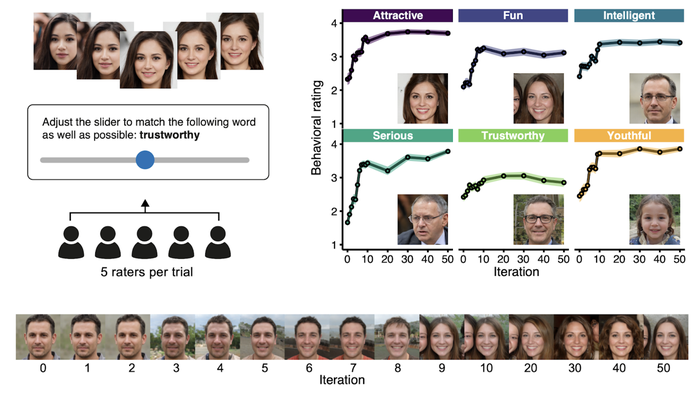

- Our face experiment connects GSP to a state-of-the-art image synthesis model, StyleGAN. Sample quality converges surprisingly early, producing faces that participants reliably identify with the chosen adjectives, e.g. 'trustworthiness'.

Conclusions

- GSP can be a powerful tool for jointly characterizing semantic representations in humans and machines.
- We reformulated MCMCP and GSP using utility theory and introduced an aggregation parameter.
- We validate the sucess of GSP with 16 experiments and 2960 participants.

Contact
rajamarjieh@gmail.com
raja.marjieh@ae.mpg.de

[1] A. Sanborn and T. L. Griffiths, "Markov Chain Monte Carlo with People," in Advances in Neural Information Processing Systems, pp. 1265–1272, 2008. [2] Tenenbaum, Joshua B., et al. "How to grow a mind: Statistics, structure, and abstraction." *science* 331.6022 (2011): 1279-1285. [3] Saxe, Andrew M., James L. McClelland, and Surya Ganguli. "A mathematical theory of semantic development in deep neural networks." PNAS 116.23 (2019): 11537-11546. [4] Dasgupta, Ishita, et al. "A theory of learning to infer." *Psychological Review* 127.3 (2020): 412. [5] E. Härkönen, A. Hertzmann, J. Lehtinen, and S. Paris, "GANSpace: Discovering interpretable GAN controls," arXiv, 2020. [6] T. L. Griffiths, D. Daniels, J. L. Austerweil, and J. B. Tenenbaum, "Subjective randomness as statistical inference," Cognitive psychology, vol. 103, pp. 85–109, 2018 [7] K. E. Train, Discrete choice methods with simulation. Cambridge University Press, 2009. [8] https://github.com/Dallinger/Dallinger