

Towards Human-Agent Communication via the Information Bottleneck Principle

Mycal Tucker
mycal@mit.edu

Julie Shah
julie_a_shah@csail.mit.edu

Roger Levy
rplevy@mit.edu

Noga Zaslavsky
nogazs@mit.edu

Abstract—Emergent communication research often focuses on optimizing task-specific utility as a driver for communication. However, human languages appear to evolve under pressure to efficiently compress meanings into communication signals by optimizing the Information Bottleneck tradeoff between informativeness and complexity. In this work, we study how trading off these three factors — utility, informativeness, and complexity — shapes emergent communication, including compared to human communication. To this end, we propose Vector-Quantized Variational Information Bottleneck (VQ-VIB), a method for training neural agents to compress inputs into discrete signals embedded in a continuous space. We train agents via VQ-VIB and compare their performance to previously proposed neural architectures in grounded environments and in a Lewis reference game. Across all neural architectures and settings, taking into account communicative informativeness benefits communication convergence rates, and penalizing communicative complexity leads to human-like lexicon sizes while maintaining high utility. Additionally, we find that VQ-VIB outperforms other discrete communication methods. This work demonstrates how fundamental principles that are believed to characterize human language evolution may inform emergent communication in artificial agents.

I. INTRODUCTION

Good communication is a critical component of successful teams of humans and artificial agents, but differing notions of “good” makes training agents to develop such communication challenging. One view of communication focuses on *utility* by framing languages as successful to the extent they enable high task performance. Emergent communication literature, wherein agents learn to communicate while optimizing a task-specific reward or utility, emphasizes this view [1–6]. However, this approach may lead to communication that is too complex for humans to understand [7] or has slow convergence rates [8], implying that additional constraints are needed in order to guide agents toward human-like communication.


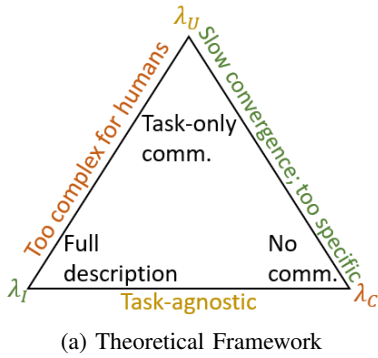
Here, we use constraints from cognitive science by building upon a recent line of work that argues that human languages evolve under pressure to efficiently compress meanings into communicative signals [9–13]. This notion of efficiency is formulated in terms of the Information Bottleneck (IB) principle [14], which is a general information-theoretic principle with broad scope in machine learning [15–19], and in this context, it can be interpreted as a tradeoff between the complexity and informativeness of communication [9]. Intuitively, complexity corresponds to the number of bits that are needed for communication, and informativeness corresponds to how well a listener can infer the speaker’s intended meaning

regardless of a specific task (e.g., humans understand what “blue” means regardless of any specific task, which may be “avoid the blue box” or “find a blue cup”). We hypothesize that taking into account complexity and informativeness, in addition to optimizing a task-specific utility function, may help guide artificial agents toward more natural, human-like communication systems.

We test this idea by training teams of artificial agents to optimize a tradeoff between maximizing task utility, maximizing communicative informativeness, and minimizing communicative complexity, as depicted in Figure 1. To this end, we extend prior work [20, 21] and propose Vector-Quantized Variational Information Bottleneck (VQ-VIB), a method for training neural agents to compress inputs into signals, while maintaining a discrete (symbolic) signal representation that is embedded in a continuous space. In experiments, we find that training to increase informativeness improves the convergence rate of emergent communication to high-performance communication, while annealing the complexity loss generates a spectrum of learned communication, which we show aligns with a variety of human languages. Additionally, we find that VQ-VIB appears better able to achieve high utility for the same complexity and informativeness than other discrete communication methods. This suggests a promising avenue for supporting human-robot communication by instilling inductive biases inspired by cognitive science.

II. RELATED WORK

Here, we describe related work in emergent communication, including in a specific color domain which we use; we include approaches for measuring informativeness and complexity in Section III. Generally, in emergent communication settings, researchers train agents to accomplish a cooperative task in which communication should enable better task performance [3, 5, 22–24]. Some works simplify communication by discretizing communication, by adding noise to the environment, or passing communication through a centralized message processor; these methods all correspond to limiting complexity [21, 25, 26]. At the same time, other works find that auxiliary losses designed to create “meaningful” communication improve convergence rates [8, 27, 28], often inspired by notions adjacent to informativeness. Color reference games have been widely used in emergent communication literature [1, 29–31], mirroring the centrality of color naming for studying humans’ semantic categorization and language



$$\text{Informativeness} := -\mathbb{E}[d(X, \hat{X})]$$

$$\text{Complexity} := I(X, C)$$

$$\text{Utility} := U(X, Y)$$

(b) Simplified Architecture. A speaker, S , maps an input, x , to a communication signal, c . A listener, L , takes an action, y , and a decoder, D , reconstructs x . Informativeness is negatively related to the expected distortion between x and \hat{x} ; complexity is defined as the mutual information between X and C ; utility can be any given task-specific reward function or loss.

Fig. 1: Our theoretical framework and training losses allow us to weigh tradeoffs between utility (λ_U), informativeness (λ_I), and complexity (λ_C) to guide properties of emergent communication.

evolution [32]. Most relevant to our work is Zaslavsky et al. [9]’s finding that color naming systems across languages (from a large-scale, cross-linguistic dataset, the Word Color Survey [WCS; 33]) attain complexity–informativeness tradeoffs that are near-optimal in the Information Bottleneck [IB; 14] sense. Chaabouni et al. [31] further showed that AI agents can learn IB-efficient color naming systems by playing color reference games (and see [30] for a related study using different notions of efficiency). While emergent communication literature has begun to rediscover the importance of complexity and informativeness in communication, these terms have rarely been explicitly recognized, whereas we include terms for utility, informativeness, and complexity in training agents.

III. TECHNICAL APPROACH

Here, we show how to adapt standard neural architectures to enable informativeness- and complexity-based objectives. We then introduce a new neural architecture, Vector-Quantized Variational Information Bottleneck (VQ-VIB), for learning complexity-limited discrete representations in a continuous space. Lastly, we combine utility, informativeness, and complexity terms into an overall objective.

A. Information Bottleneck for Semantic Communication

We adopt and extend Zaslavsky et al. [9]’s widely-supported IB framework for semantic communication [10–13]. In this framework, speakers and listeners jointly optimize the IB tradeoff between the complexity and informativeness of their communication system. A speaker is characterized as a probabilistic encoder $q(c|x)$ that, given an input, $x \sim p(x)$, generates a signal, c . A listener is characterized as a probabilistic decoder $q(\hat{x}|c)$, that given c aims to reconstruct the speaker’s input. Complexity is measured by the mutual information between the speaker’s inputs and signals, $I(X, C)$. Informativeness is measured by the degree to which the speaker’s mental representation of x matches the listener’s mental representation of \hat{x} . In general, these mental representations are defined by probability distributions, m and \hat{m} respectively, over some feature space. In that case, minimizing the expected Kullback-Leibler (KL) divergence $\mathbb{E}[D[m||\hat{m}]]$ amounts to maximizing the informativeness of the communication. This yields the

following optimization principle: $\min I(X, C) - \beta \mathbb{E}[D[m||\hat{m}]]$, where β is a tradeoff parameter. It has been shown that this principle is equivalent to the IB principle [9].

Solving this optimization problem directly is challenging in large scale domains and with agents that are parameterized by neural networks. To address this, we take a Variational Information bottleneck [VIB; 18] approach, in which an upper bound is optimized, similar to ELBO in Variational Autoencoders [VAEs, 34] and β -VAEs [35]. We also approximate informativeness by an autoencoding loss, $d(x, \hat{x})$, defined by the mean squared error (MSE) loss of the speaker’s input and the listener’s reconstruction.

B. Vector Quantized Variational Information Bottleneck

Our proposed neural architecture, the Vector Quantized Variational Information Bottleneck (VQ-VIB), learns information-constrained discrete communication in a continuous space. Our technique combines prior work in discrete representation learning [VQ-VAE, 20] with β -VAE [35].

1) *Vector-Quantized Variational Autoencoder*: First, consider a simplified Vector-Quantized Variational Autoencoder (VQ-VAE) architecture [20]. As in standard VAEs, an encoder is a conditional distribution over latent representations, given an input: $q(z|x)$. A deterministic network maps from an input, x , to a latent, z : we dub such encodings, $z(x)$. Next, for a VQ-VAE parametrized with K embedding vectors, ζ , in the latent space, $z(x)$ is discretized to a particular embedding vector, ζ_i , by mapping to the nearest embedding vector.

2) *β Variational Autoencoder (β -VAE)*: We also use notions from β -VAE, a variational approach for generating information-constrained continuous representations with neural nets. Following Higgins et al. [35], an encoder network maps from an input to μ and σ , which are used as parameters to a Gaussian distribution from which a latent representation, z , is sampled. β -VAE networks are trained via a weighted combination of a reconstruction loss and (β times) the KL divergence of $\mu(x)$ and $\sigma(x)$ from a unit Gaussian; this KL loss corresponds to an upper bound of $I(X, Z)$.

3) *VQ-VIB Architecture*: Here, we integrate VQ-VAE and β -VAE for a method of generating discrete, information-bounded representations. Our architecture, Vector-Quantized




Fig. 2: VQ-VIB architecture blends Vector Quantization and Variational Information Bottleneck. The speaker encodes an input, x , as parameters to a Gaussian, from which a communication vector is sampled and then discretized by mapping to the nearest encoding, to produce communication, c .




Fig. 3: In particle-world environments, a speaker (clear) communicated to a listener (gray) to reach a desired landmark.

Variational Information Bottleneck (VQ-VIB), corresponds to first using an encoder to sample a representation z in a continuous space, followed by a quantization layer that discretizes the sample by mapping to the nearest encoding ζ . This is depicted in Figure 2 and is formally defined in Equation 1. The likelihood of a communication vector, c , given an input, x , corresponds to the likelihood that an encoding z , sampled with parameters $\mu(x)$ and $\sigma(x)$, is closest to an embedding vector equal to c (Eq. 1). This intuitively corresponds to notions of uncertainty over which “word” to use. One may compute an upper bound on $I(X, C)$ based on this uncertainty by using the standard KL divergence used for Gaussians in a β -VAE, using $\mu(x)$ and $\sigma(x)$.

$$\begin{aligned} z &\sim \mathcal{N}(\mu(x), \sigma(x)) \\ q(c|x) &= \mathbb{P}(c = \operatorname{argmin}_{\zeta} [(z - \zeta)^2]) \end{aligned} \quad (1)$$

C. Overall Optimization

Here, we combine losses into an idealized objective function to maximize, shown in Equation 2. This objective function is the weighted combination of three terms, wherein we seek to maximize a tradeoff between high utility, high informativeness (low expected distortion d), and low complexity. Intuitively, these three scalar weights form a simplex that describe a range of communication from silence (minimizing complexity) to task-optimized communication that does not generalize (maximizing utility) to complete descriptions (maximizing informativeness).

$$\text{maximize } \lambda_U U(X, Y) - \lambda_I \mathbb{E}[d(X, \hat{X})] - \lambda_C I(X, C) \quad (2)$$




Fig. 4: In the color reference game, agents observed colors from the WCS Dataset [33].




Fig. 5: **Top:** In the 9 points environment, greater λ_I led to faster convergence for VQ-VIB; **Bottom:** Annealing λ_C decreased reward and complexity for all methods.

In practice, we train neural agents via gradient descent to maximize approximations of these terms. Utility is represented via the policy gradient (for RL settings) or a supervised loss. Informativeness is computed via the negative MSE between the speaker’s input and a decoder’s output, given the communication. Lastly, complexity is bounded by the KL divergence between the speaker’s distribution over communication and a fixed marginal (either a uniform distribution for onehot communication, or a unit Gaussian for others).

IV. RESULTS

A. Particle Worlds

We trained agents in two types of 2D particle worlds, as depicted in Figure 3. In each environment, the speaker observed a “target” landmark (in 9 points, there were 9 landmarks at fixed locations, whereas in Uniform, the landmarks were spawned a random locations) and communicated to the listener. The team reward was the negative distance from the listener to the target, so agents learned to communicate about the target location. We trained agents with Multi-Agent Deep Deterministic Policy Gradient [MADDPG; 3] and compared continuous (real-valued), onehot, VQ-VIB, and prototype-based [Proto; 21] communication. All except continuous communication are fundamentally discrete. For all environments and architectures,




Fig. 6: Color reference game results. Left: By annealing complexity, we recreated a range of near-optimal color-naming schemes (black curve is the theoretical IB bound from [9]). Middle: 2D PCA of VQ-VIB Before communication shows relationships between different clusters in the communication space. Right: mode maps of the agent (no box) and human (dashed box) color naming systems, where each color chip from Figure 4 is colored by the centroid of its modal category.

training agents with higher λ_I improved training convergence, as depicted for VQ-VIB in the 9-points environment in Fig. 5 (top). Furthermore, as we increased λ_C , all architectures learned less complex, but also lower-utility, communication (Fig. 5 bottom). However, VQ-VIB achieved higher utility, for the same complexity, than other discrete communication methods. These trends also appeared in the Uniform environment: 1) higher λ_I improved convergence, 2) increasing λ_C created a spectrum of communication complexity and 3) VQ-VIB outperformed other discrete methods.

B. Color Reference Games

We also trained agents in a color reference game in which a speaker agent observed a target color (drawn from WCS data, as shown in Fig. 4), and a listener agent had to distinguish between the target and a distractor color. As in the previous experiments, we found that increasing λ_I , for a fixed λ_C , improved convergence rates, and VQ-VIB agents typically learned more complex communication than onehot agents.

We tested agents for $\lambda_I = 1.0$ while linearly annealing λ_C to penalize complexity. As λ_C increased, the communication complexity decreased (Figure 6 left). Over 5 random seeds, the points in the plot, generated at each training epoch, closely approximated the IB upper bound for maximum informativeness for a given complexity. The learned color naming schemes at different complexity levels exhibited important patterns (Figure 6 right). As complexity decreased, the number of communication clusters emitted by VQ-VIB agents decreased (visualized via 2-dimensional principle component analysis). Clustering in the communication space was reflected in the color space (e.g., “yellow” and “orange” merging). More generally, a least-squares linear correlation between the distance between pairs of communication vectors and the distance between associated colors (in perceptually-based CIELAB space) showed $r^2 = 0.97 \pm (0.01)$ for VQ-VIB, whereas for onehot

communication $r^2 = 0.76 \pm (0.03)$. That is, for VQ-VIB, if two communication vectors were far apart in communication space, they described colors that were far apart. This shows a fundamental connection between the form and meaning of communication used by VQ-VIB agents, which onehot-based communication is architecturally unable to support.

Lastly, we compared the agent color naming schemes to human language data. For each of the 110 languages in the WCS Dataset [33], we found the best-matching agent as determined by gNID, a measure of dissimilarity of naming distributions introduced by Zaslavsky et al. [9]. Example nearest languages are shown in Figure 6 in dashed boxes. Annealing complexity of artificial communication was necessary to generate similar naming conventions for each human language. Furthermore, while onehot and VQ-VIB Before achieved similar informativeness for the same complexity, VQ-VIB Before reached higher utility, indicating an architectural advantage.

V. CONCLUSION

In this work, we explore the interplay of utility, informativeness, and complexity in emergent communication. To the best of our knowledge, we are the first to directly optimize agents to communicate according to a combination of all three terms. Rewarding informativeness led to faster communication convergence, and penalizing complexity created more human-like communication. While we focused on training agents entirely in “self-play,” natural extensions of this work would include supporting human-robot communication, wherein robots could learn partner-specific communication by appropriately constraining complexity and a small amount of supervised data.

ACKNOWLEDGMENTS

M.T. is supported by an Amazon Science Hub Fellowship. N.Z. is supported by a K. Lisa Yang Integrative Computational Neuroscience (ICoN) Postdoctoral Fellowship.

REFERENCES

- [1] Luc Steels and Tony Belpaeme. Coordinating perceptually grounded categories through language: A case study for colour. *Behavioral and Brain Sciences*, 28(4):469–489, 2005.
- [2] Jakob Foerster, Ioannis Alexandros Assael, Nando de Freitas, and Shimon Whiteson. Learning to communicate with deep multi-agent reinforcement learning. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. URL <https://proceedings.neurips.cc/paper/2016/file/c7635bfd99248a2cdef8249ef7fbef4-Paper.pdf>.
- [3] Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. *Neural Information Processing Systems (NIPS)*, 2017.
- [4] Igor Mordatch and Pieter Abbeel. Emergence of grounded compositional language in multi-agent populations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [5] Angeliki Lazaridou, Alexander Peysakhovich, and Marco Baroni. Multi-agent cooperation and the emergence of (natural) language. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL <https://openreview.net/forum?id=Hk8N3ScIlg>.
- [6] Angeliki Lazaridou and Marco Baroni. Emergent multi-agent communication in the deep learning era. *arXiv preprint arXiv:2006.02419*, 2020.
- [7] Rahma Chaabouni, Eugene Kharitonov, Emmanuel Dupoux, and Marco Baroni. Anti-efficient encoding in emergent communication. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/31ca0ca71184bbdb3de7b20a51e88e90-Paper.pdf>.
- [8] Tom Eccles, Yoram Bachrach, Guy Lever, Angeliki Lazaridou, and Thore Graepel. Biases for emergent communication in multi-agent reinforcement learning. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/fe5e7cb609bde6d62449d61849c38b0-Paper.pdf>.
- [9] Noga Zaslavsky, Charles Kemp, Terry Regier, and Naftali Tishby. Efficient compression in color naming and its evolution. *Proceedings of the National Academy of Sciences*, 115(31):7937–7942, 2018. doi: 10.1073/pnas.1800521115. URL <https://www.pnas.org/doi/abs/10.1073/pnas.1800521115>.
- [10] Noga Zaslavsky, Terry Regier, Naftali Tishby, and Charles Kemp. Semantic categories of artifacts and animals reflect efficient coding. In *Proceedings of the 41st Annual Meeting of the Cognitive Science Society*, 2019.
- [11] Noga Zaslavsky, Mora Maldonado, and Jennifer Culbertson. Let’s talk (efficiently) about us: Person systems achieve near-optimal compression. In *Proceedings of the 43rd Annual Meeting of the Cognitive Science Society*, 2021.
- [12] Noga Zaslavsky, Karee Garvin, Charles Kemp, Naftali Tishby, and Terry Regier. The evolution of color naming reflects pressure for efficiency: Evidence from the recent past. *Journal of Language Evolution*, 2022.
- [13] Francis Mollica, Geoff Bacon, Noga Zaslavsky, Yang Xu, Terry Regier, and Charles Kemp. The forms and meanings of grammatical markers support efficient communication. *Proceedings of the National Academy of Sciences*, 118(49), 2021.
- [14] Naftali Tishby, Fernando C. Pereira, and William Bialek. The Information Bottleneck method. In *Proceedings of the 37th Annual Allerton Conference on Communication, Control and Computing*, 1999.
- [15] Gal Chechik, Amir Globerson, Naftali Tishby, and Yair Weiss. Information Bottleneck for Gaussian variables. *Journal of Machine Learning Research*, 6, 2005.
- [16] Ohad Shamir, Sivan Sabato, and Naftali Tishby. Learning and generalization with the Information Bottleneck. *Theoretical Computer Science*, 411(29-30):2696–2711, 2010.
- [17] Naftali Tishby and Noga Zaslavsky. Deep learning and the information bottleneck principle. In *2015 IEEE Information Theory Workshop (ITW)*, pages 1–5, 2015. doi: 10.1109/ITW.2015.7133169.
- [18] Alex Alemi, Ian Fischer, Josh Dillon, and Kevin Murphy. Deep variational information bottleneck. In *ICLR*, 2017. URL <https://arxiv.org/abs/1612.00410>.
- [19] Ravid Shwartz-Ziv and Naftali Tishby. Opening the black box of deep neural networks via information. *arXiv preprint arXiv:1703.00810*, 2017.
- [20] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017.
- [21] Mycal Tucker, Huao Li, Siddharth Agrawal, Dana Hughes, Katia P. Sycara, Michael Lewis, and Julie Shah. Emergent discrete communication in semantic spaces. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=EuvW3-Lmoxm>.
- [22] Kalesha Bullard, Franziska Meier, Douwe Kiela, Joelle Pineau, and Jakob Foerster. Exploring zero-shot emergent communication in embodied multi-agent populations. *arXiv preprint arXiv:2010.15896*, 2020.
- [23] Kalesha Bullard, Douwe Kiela, Joelle Pineau, and Jakob Foerster. Quasi-equivalence discovery for zero-shot emergent communication. *arXiv preprint*

- arXiv:2103.08067*, 2021.
- [24] Angeliki Lazaridou, Karl Moritz Hermann, Karl Tuyls, and Stephen Clark. Emergence of linguistic communication from referential games with symbolic and pixel input. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. URL <https://openreview.net/forum?id=HJGv1Z-AW>.
 - [25] Łukasz Kuciński, Tomasz Korbak, Paweł Kołodziej, and Piotr Miłoś. Catalytic role of noise and necessity of inductive biases in the emergence of compositional communication. *Advances in Neural Information Processing Systems*, 34:23075–23088, 2021.
 - [26] Rundong Wang, Xu He, Runsheng Yu, Wei Qiu, Bo An, and Zinovi Rabinovich. Learning efficient multi-agent communication: An information bottleneck approach. In *ICML*, pages 9908–9918, 2020. URL <http://proceedings.mlr.press/v119/wang20i.html>.
 - [27] Natasha Jaques, Angeliki Lazaridou, Edward Hughes, Caglar Gulcehre, Pedro Ortega, DJ Strouse, Joel Z Leibo, and Nando De Freitas. Social influence as intrinsic motivation for multi-agent deep reinforcement learning. In *International Conference on Machine Learning*, pages 3040–3049. PMLR, 2019.
 - [28] Toru Lin, Jacob Huh, Christopher Stauffer, Ser Nam Lim, and Phillip Isola. Learning to ground multi-agent communication with autoencoders. *Advances in Neural Information Processing Systems*, 34, 2021.
 - [29] Mike Dowman. Explaining color term typology with an evolutionary model. *Cognitive Science*, 31(1):99–132, 2007.
 - [30] Mikael Kågebäck, Emil Carlsson, Devdatt Dubhashi, and Asad Sayeed. A reinforcement-learning approach to efficient communication. *PLOS ONE*, 15(7):1–26, 07 2020.
 - [31] Rahma Chaabouni, Eugene Kharitonov, Emmanuel Dupoux, and Marco Baroni. Communicating artificial neural networks develop efficient color-naming systems. *Proceedings of the National Academy of Sciences*, 118(12):e2016569118, 2021. doi: 10.1073/pnas.2016569118. URL <https://www.pnas.org/doi/abs/10.1073/pnas.2016569118>.
 - [32] Brent Berlin and Paul Kay. *Basic Color Terms: Their Universality and Evolution*. University of California Press, Berkeley and Los Angeles, 1969.
 - [33] Paul Kay, Brent Berlin, Luisa Maffi, William R. Merrifield, and Richard Cook. *The World Color Survey*. Stanford: Center for the Study of Language and Information, 2009.
 - [34] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
 - [35] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational

framework. 2016.