

Decentralised Emergence of Robust and Adaptive Linguistic Conventions in Populations of Autonomous Agents Grounded in Continuous Worlds: Supplementary Material

Anonymous Author(s)

1 CODE AND DATA

The code can be found within the ‘code’ zip folder. It contains all the necessary resources to replicate and run the experiments described in the paper. Inside, you’ll find three main components: ‘data’, ‘exp’, and ‘plot’. The ‘data’ folder houses the code to download, preprocess, and format the datasets. The ‘exp’ folder houses the code to reproduce the experiments of the paper. The ‘plot’ folder houses the code to reproduce the figures of the paper.

2 ENVIRONMENT SETUP

This section details the process by which the environments of the experiments are setup. Training and test splits for the three datasets were created in a two-stage process. The first stage consisted in the creation of scenes, i.e. unique sets of entities that can serve as the context for a language game (see the *Context selection* step in Section 3 of the paper). In other words, a scene represents the ‘environment’ in which the communication interaction takes place, i.e. the shared context. Each scene consists of 3 to 10 entities, with entities occurring in a training scene being excluded from being part of a test scene. For CLEVR, the distribution of scenes from the original training and test splits were kept, holding 70,000 and 15,000 scenes respectively. In the case of WINE, 90% of the wine samples were used to create 20,000 training scenes, and the remaining 10% were used to create 1,000 test scenes. For CREDIT, 90% of the financial transactions were used to create 40,000 training scenes and the remaining 10% were used to create 4,000 test scenes. In a second stage, the actual training and tests sets were constructed by randomly sampling from the training and test scenes. For each dataset, training and test sets consisting of 1,000,000 and 100,000 scenes were compiled. The same scene can thus occur multiple times in the training or test set, but can never be part of both. The fact that the same scenes can occur multiple times does not entail that the same game is played multiple times. Indeed, many different games can be played in the same scene depending on the participating agents and the selected topic.

3 HARDWARE, TRAINING REGIME, TUNED HYPERPARAMETERS

All experiments were conducted on a 20-core INTEL Xeon Gold 6148 processor, paired with 32GB of RAM. One million sequential games (the amount of games in each experiment) were executed on this hardware in ± 8 hours. Table 1 includes the space of hyperparameters explored for the baseline CLEVR experiment. The best performing set of hyperparameters (in terms of communicative success and lexicon coherence) are reported in the main text. Every subsequent experiment uses this same set of parameters. The parameters do not seem to be domain dependent as the best

Table 1: Overview of hyperparameter search

Param.	Tested values
s_r	$\{+0.01, +0.05, +0.1\}$
s_p	$\{-0.01, -0.05, -0.1\}$
s_{li}	$\{-0.05, -0.01, -0.02, -0.05, -0.1\} * \text{sim}(c_q, c_r)$
σ_i	$\{0.001, 0.005, 0.01, 0.05, 0.1\}$
ω_i	$\{0.1, 0.2, 0.5, 0.75, 1.0\}$
c_r	$\{+1, +5, +10\}$
c_p	$\{-1, -5, -10\}$

performing set for CLEVR generalises very well on the other two datasets.