Interaction history as a source of compositionality in emergent communication

Tomasz Korbak,[1,2] Julian Zubek,[2] Łukasz Kuciński,[3] Piotr Miłoś,[3,4] and Joanna Rączaszek-Leonardi[2]
[1] University of Sussex | [2] University of Warsaw | [3] Polish Academy of Sciences | [4] University of Oxford

In this paper, we explore interaction history as a particular source of pressure for achieving emergent compositional communication in multi-agent systems. We propose a training regime implementing template transfer, the idea of carrying over learned biases across contexts. In the presented method, a sender-receiver dyad is first trained with a disentangled pair of objectives, and then the receiver is transferred to train a new sender with a standard objective. Unlike other methods (e.g. the obverter algorithm), the template transfer approach does not require imposing inductive biases on the architecture of the agents. We experimentally show the emergence of compositional communication using topographical similarity, zero-shot generalization and context-independence as evaluation metrics. The presented approach is connected to an important line of work in semiotics and developmental psycholinguistics: it supports a conjecture that compositional communication is scaffolded on simpler communication protocols.

Keywords: emergent communication, path-dependence, compositionality, Lewis signaling games, generalization

1. Introduction

Language-like communication protocols can emerge in games that require the agents to share information and coordinate behaviour (Foerster et al., 2016; Lazaridou et al., 2016; Jaques et al., 2018; Chaabouni et al., 2020; Kharitonov and Baroni, 2020). One important feature of human languages is *compositionality* – there are complex signals constructed through the combination of signals. Compositionality is considered a key feature of general intelligence because it

https://doi.org/10.1075/is.21020.kor
Interaction Studies 22:2 (2021), pp. 212–243. ISSN 1572-0373 | E-ISSN 1572-0381

facilitates generalization (adaptability to novel situations) and productivity (an infinite number of meanings can be created using a finite set of primitives) (Lake et al., 2016). However, recent work on emergent languages in artificial intelligence shows that under certain circumstances compositionality is hard to achieve and requires strong inductive biases to be imposed on the agents (Kottur et al., 2017).

We recognize that the evolutionary pressure for compositional communication may come from multiple sources: inductive biases, structure of the environment, social situation, learnability, etc. In this paper we isolate and model one particular source of such pressure connected with the history of interactions in a changing context. We demonstrate that communication protocols exhibiting compositionality can emerge via adaptation of pre-existing, simpler non-compositional protocols to a new context. This procedure is an instance of *template transfer* (Barrett and Skyrms, 2017). Our model implements the idea of template transfer by sharing agents across signaling games of varying complexity. We decompose learning compositional communication into three phases: (i) learning a visual classifier, (ii) learning non-compositional communication protocols, and (iii) learning a compositional communication protocol. This decomposition closely follows distinctions established in semiotics (the hierarchy of (i) icons, (ii) indices, and (iii) symbols postulated by Peirce (1998), see Section 5) and is more plausible in the light of human language development than other approaches. Crucially, the biases learned in simple games in phase (ii) are sufficient to incentivize a compositional communication protocol to emerge in phase (iii). The incentive for compositional communication does not come from their innate inductive biases but from the history of their involvement in different games. We compare the template transfer approach with other method of achieving compositionality – the obverter algorithm (Batali, 1998; Choi et al., 2018)–on three different metrics: zero-shot generalization, context-independence and topographical similarity. The results demonstrate that the ability to communicate compositionality can emerge in a model less cognitively demanding than the obverter approach.

2. Related work

In this paper, we assume a broadly pragmatic or game-theoretic approach to language evolution: communication emerges as a tool for guiding joint action or for enabling coordination in a multi-agent system trained with a joint objective

(Raczaszek-Leonardi et al., 2018).[1] Lewis signaling games (Lewis, 1969; Skyrms, 2010) are a popular game-theoretic model of communication. A Lewis signaling games consists of a set of states $x \in X$, a set of available actions $y \in Y$, a set of available messages $m \in M$, and a sender s (sending a message m upon observing a state x), and receiver r (taking an action y upon receiving a message m), and finally a loss function $\mathcal{L}(x, y)$ assigning each state-action pair a reward or penalty. With a particular choice of $\mathcal{L}$, communication can be modeled as having the sender and the receiver agreeing upon a communication protocol that allows the receiver to take advantage of the information about the state x available (only) to the sender to take the action y minimizing the loss function $\mathcal{L}(x, y)$.

Computational models of Lewis signaling games traditionally relied on either simple reinforcement learning (e.g. Roth-Erev model; Skyrms (2010)) or evolutionary optimization (Cangelosi, 2001; Grouchy et al., 2016) for learning the parameters θ and Ψ of the policies for the sender s_θ and the receiver r_Ψ. With the rise of deep learning (Goodfellow et al., 2016), deep neural networks started being used to implement policies of the agents with parameters optimized via gradient descent implemented using the backpropagation algorithm (Rumelhart et al., 1986). In a typical setting, learning boils down to descending along the gradient $\nabla_{\theta,\Psi} \mathcal{L}(r_\Psi(s_\theta(x)),y)$ for a loss function $\mathcal{L}$, a state x and an action y. The introduction of more powerful models (in terms of capacity) and more efficient training regimes (in terms of convergence time) contributed to the emergence of qualitatively novel phenomena (e.g. counter-factual reasoning (Jaques et al., 2018)) as well as enabled using more psychologically realistic settings (e.g. presenting the agents with raw visual inputs as opposed to pre-processed, discrete representations of stimuli (Lazaridou et al., 2018; Bouchacourt and Baroni, 2018).

Inductive biases for compositional communication

Kottur et al. (2017) argue that the emergence of compositionality requires strong inductive biases to be imposed on communicating agents. In a guessing game with inputs being objects characterized by color and shape, agents implemented by a vanilla architecture (i.e. without additional constrains motivated by compositionality) will most likely end up developing an information-theoretically optimal yet non-compositional communication protocol – a hash function for the objects – that will show poor generalization to novel combinations of colors and shapes (Kottur et al., 2017). One recurring approach to enforce compositionality is plac-

1. There are also more syntactically or semantically oriented approaches, which focus more on the development of lexicons and grammars in single agents. For a broader review, see (Gong et al., 2014).

ing pressure on agents to use symbols consistently across varying contexts. To that end, Kottur et al. (2017) and Das et al. (2017) reset the memory of an agent between producing or receiving subsequent parts of a message, which helps to obtain a consistent symbol grounding (i.e. each symbol is associated with a shape irrespective of color or with color irrespective of shape). Resetting the memory of an agent in the middle of receiving or producing a message can be argued to be an *ad hoc* manipulation, however, which is of limited interest to researchers focused on uncovering biologically plausible mechanisms of compositionality.

Obverter approach

A more psychologically plausible approach is explored by Choi et al. (2018) and Bogin et al. (2018), who take inspiration from the obverter algorithm (Oliphant and Batali, 1997; Batali, 1998). The obverter (from the Latin *obverto,* to turn towards) algorithm is based on the assumption that an agent can use its own responses to messages to predict other agent's responses, and thus can iteratively compose its messages to maximize the probability of the desired response. In a typical game, two agents a_ξ and a_k (with policies parametrized by ξ and k) exchange the roles of the sender and the receiver. If an agent is the receiver ($a_\xi = r_\xi$), it behaves as in the object naming game. If an agent is the sender ($a_\xi = s_\xi$), it sends message that would have produced the optimal response (to the best of a_ξ 's knowledge) if a_ξ had received such a message as a receiver. More formally, a_ξ sends a message $m = \arg\max_{m'} r_\xi(\mathrm{y}_{c|}\mathrm{m}')$, assuming a correct action y_c is known or can be predicted by a_ξ. This can be interpreted as agents possessing a *theory of mind* (Bruner, 1981; Tomasello et al., 2005) or a model for predicting the response of the other agent $r_K(y|m)$ based on own their policy $r_\xi(y|m)$.

A limitation of the obverter is that it makes strong assumptions about the agents and task: to be able to use themselves as models of others, the agents should share a similar architecture[2] and the task must be symmetric (the agents must be able to exchange their roles). This excludes games with functional specialization of agents.

Another problem is the computational complexity of the decoding procedure. Even assuming greedy decoding (i.e. that the sender will compose a message by progressively choosing the next symbol m_t maximizing $r\xi(y_c|m_{1:t})$), producing a

2. While Batali (1998) and Choi et al. (2018) used the same architecture of the agents, one can imagine relaxing this assumption and only using *similar* architecture. However, it no longer guarantees that the sender's model class (the family of probability distributions over symbol sequences it can represent) is rich enough to represent the distribution of the receiver.

message requires O($|V|T$) queries to the model of the receiver (where $|V|$ is vocabulary size and T is maximum message length).

Population-based training

A different family of approaches tries to incetivize compositionality by training entire populations of senders and receivers and creating a pressure for learnability of the communication protocol for new agents. This approach was pioneered by the iterated learning model, which assumed that agents acquire a communication protocol by being (implicitly or explicitly) taught by the agents from previous generations (Kirby, 2001). The cultural transmission is imperfect, which creates a bias towards protocols that are both expressive and easy to teach (Brighton, 2002). Iterated learning was found to lead to compositionality both in computational experiments (Brighton, 2002) as well as in experiments with human subjects (Kirby et al., 2008). In the machine learning literature, generational transmission as a mechanism for inducing compositionality was explored by Li and Bowling (2019) and Cogswell et al. (2019), who simulate the arrival of new language users by periodically resetting weights of some agents in the population. Their experiments corroborated the effect of increased compositionality and found it to be complementary with other factors that encourage compositionality. Injecting noise into the messages produces a similar increase in compositionality (Kuciński et al., 2021).

Multi-task training

Yet another approach, most similar in spirit to ours, was introduced by De Beule and Bergen (2006). In this work, a population of agents plays a guessing game in a world populated by events involving agents and patients. There are N_e event predicates (e.g. *kicked)* and N_p person predicates (e.g. *Mary),* giving rise to $2N_pN_e$ structured topics and $N_p + N_e$ atomic topics. The fraction between the number of structured topics presented to the agents and the number of atomic topics presented to the agents is known as *task complexity.* Task complexity turns out to be a crucial parameter influencing compositionality. For an intuitive explanation, consider the event "Mary loves Eve". A sender who has never seen neither event predicate *love* nor person predicates *Mary* and *Eve* might employ a new word to communicate this event. However, a sender already knowing the word for *Mary* might reuse it together with new words for novel elements of the event. The experiments conducted by De Beule and Bergen (2006) demonstrate that the incentive to reuse known symbols leads to the emergence of compositional communication in games with low yet non-zero task complexity, i.e. when agents communicate mostly about atomic topics but also about structured topics. Contributing to this line of thinking, we show how a similar effect of reusing parts of a non-compositional communication protocol in a compositional fashion can emerge when training with structured topics occurs after (not simultaneously to) training with atomic topics.

Evolutionary origins of grammar

Nowak and Krakauer (1999) presented a game-theoretic model of language evolution, in which they investigated how specific features of language provide fitness advantage under certain circumstances. They suggested that simple compositional (and grammatical) languages may evolve when communicating different aspects of the world is rewarded independently (disentangled rewards). For example, it may be evolutionary beneficial to name a type of danger (leopard), and it would be even more beneficial to name it and specify additional circumstances (a stalking leopard is more dangerous than a sleeping one). At first, a simple non-compositional language which only names the type of danger may evolve, and then it could be transformed into a compositional one conveying both the type of danger and circumstances. This is conceptualy similar to the template transfer approach since a) more complex language protocol relies on a simpler one, b) structure of the compositional protocol reflects structure of the rewards received by an agent.

Generalized signaling games

Barrett and Skyrms (2017) recently developed a theoretical framework of *generalized Lewis signaling games* for modeling how Lewis signaling games can be composed and transferred to new settings to yield more powerful Lewis signaling games. These generalizations can be understood in terms of *ritualization:* the process of exploiting pre-existing patterns of behavior of some agent a_1 by some other agent a_2 for the benefit of a_2. This notion gives rise to the following classes of generalized signaling games:

1. a cue-reading game is where $a_1 = s_\theta$ and $a_2 = r_\Psi$, i.e. θ is approximately fixed and the receiver takes advantage of the policy of the sender,
2. a sensory-manipulation game is where $a_1 = r_\Psi$ and $a_2 = s_\theta$, i.e. Ψ is approximately fixed and the sender takes advantage of the policy of the receiver,
3. a (proper) signaling game is where the sender and receiver both take advantage of each other's dispositions.

Barrett and Skyrms (2017) offer the following examples. For cue-reading, consider cross-species signaling networks such as hornbills receiving, understanding and exploiting alarm calls of Diana monkeys (Rainey et al., 2004). (The two species have common predators.) For an example of sensory-manipulation, consider mating rituals of frogs from the *Physalaemus pustulosus* species group. Here males of several species of *Physalaemus pustulosus* exploit the sensitivity of females for certain sounds that is evolutionary antecedent (pre-existing) and shared between *Physalaemus pustulosus* species (Ryan and Rand, 1993).

In case when the policy of a_1 evolved as a solution to a previous signaling game go between a_1 and some a_0, the new signaling game g_1 with a_1 and a_2 can be seen as evolving out of g_0. This appropriation of a policy of a_1 from g_0 to a new game g_1 is known as *template transfer.* The policy of a_2 can then be seen as a translating the inputs from the g_1 to inputs from g_0 or emulating g_0. This is why the transferred policy of a_1 might be successful in a context g_1 other than the one the policy initially evolved for (i.e. g_0).

A related phenomenon, *modular composition,* occurs when the output (i.e. receiver's action) of one game g_0 is the (sender's) input to a new game g_1, thus forming a composite game. For instance, an initial game g_0 with agents s_0, r_0 can be interpreted as itself being a policy of an agent $s_1 : x \mapsto r_0(s_0(x))$ who can then play with a new receiver r_1 thus forming a composite game $g_1 : x \mapsto r_1(r_0(s_0(x)))$. This instance of modular composition is known as *polymerization* and boils down to agents forming a signaling chain. Modular composition may also involve games with several senders and/or receivers and networks with branched flow of messages. Barrett and Skyrms (2017) provide an example of NAND games (i.e. games with two senders communicating with one receiver to jointly emulate a NAND gate) being composed to form an OR game (or, by extension, emulating an arbitrary Boolean function).

While transferred policies and solutions to composite games could in principle have evolved from scratch, template transfer and modular composition lead to orders of magnitude faster convergence. Moreover, they seem to implement a general principle of modular reuse in nature. It seems that a great deal of cognitive, social and semiotic phenomena can emerge through recursive modular composition or iterative template transfer from simpler to more complex games. This includes logical inference (Barrett and Skyrms, 2017), knowledge sharing in a community (Barrett et al., 2019) and functional specialization of agents (Barrett et al., 2018).

Motivation for our approach

The direction pursued in this paper is to explore an alternative, novel approach to the emergence of compositional communication, where the history of previous interactions can affect the emerging communication protocol, even if pressures shaping the communication change. It can be illustrated through a story inspired by the classic Wittgenstein's example of stonemasons building a construction out of stones (Wittgenstein, 1953, Paragraph 2). Stones have various shapes and colors. There are four stonemasons altogether: a master and three apprentices. The role of an apprentice is to pick a stone from a pile, name its properties, and pass it to the master. The master takes the stone without looking, and uses it in the construction. The master is working with apprentice A on a piece of wall which needs to be sturdy but not pretty, and with apprentice B on a piece of decorative wall which needs to be pretty but not sturdy. When communicating with apprentice A only stone shapes are important, when communicating with apprentice B only stone colors are important. The master simultaneously learns to communicate with apprentices A and B. Then, the work begins on a third piece of wall which needs to be both sturdy and pretty. The master hires apprentice C for the job, and they need to learn to communicate effectively. Apprentice C is initially clueless, but the master already has certain communicational habits from working with A and B. Apprentice C picks up on those and starts using compositional expressions for shape and color based on expressions used previously by A and B. Compositional communication emerges even though the initial reason for disentangling shape and color disappears.

Following Skyrms (2010), we will pose the problem as a *naming game*–a Lewis signaling game, where the sender (apprentice from the example) sees an object with two independent factors of variation (shape and color) and the receiver (master from the example) must, independently, indicate both of these factors. The objects will be presented to the sender as raw pixel data, which is motivated by (relative) biological plausibility of this setting. The receiver is disembodied in the sense that it does not receive any perceptual input, only sender's messages (this comes in contrast to another popular setup where the receiver is presented a target object and a set of distractors and then has to choose the target based on the messages). There will be a pre-defined communication channel: a fixed set of fixed-length messages composed from symbols from a fixed vocabulary. There will be an implicit temporal dimensions in the model as the sender produces the messages symbol-by-symbol and the receiver receives them symbol-by-symbol. Both agents will be implemented as recurrent neural networks with time-steps corresponding to subsequent symbols.

The aim is to explore solutions to the object naming game based not on injecting inductive biases into the architecture of the agents, but leveraging constraints established by the history of previous interactions in a game-theoretically principled manner. More specifically, we will investigate whether template transfer (as described above) can be employed as a way of achieving compositional communication in an object naming game. The character of this work is proof of concept. Relatively simple architecture of our agents, disembodied receiver and simplicity of the stimuli guarantee that the observed effects are due to template transfer, not other factors.

3. Method

In this section, we describe the experimental setup in more detail, derive the specific loss functions used in experiments and present the *template transfer* approach – the main contribution of the paper.

3.1 Experimental setup

3.1.1 *Dataset*

We conduct our experiments on a dataset consisting of 2500 images of colored three-dimensional objects. Each image has dimension of $128 \times 128 \times 3$ pixels. The dataset includes images of five shapes (box, sphere, cylinder, torus, ellipsoid) and five colors (blue, cyan, gray, green, magenta). One hundred images generated using POV-Ray ray tracing engine,[3] differing in the position of the object on a surface, are included for each color-shape pair. An analogous dataset was previously used by Lazaridou et al. (2018); Choi et al. (2018) and Bogin et al. (2018). We choose pairs for the test set by taking one of each figure and color, i.e. the test set is composed of blue boxes, cyan cylinders, gray ellipsoids, green spheres and magenta tori. Example images from the training set are shown on Figure 1.

3. The dataset was generated using code available from https://github.com/benbogin/obverter.

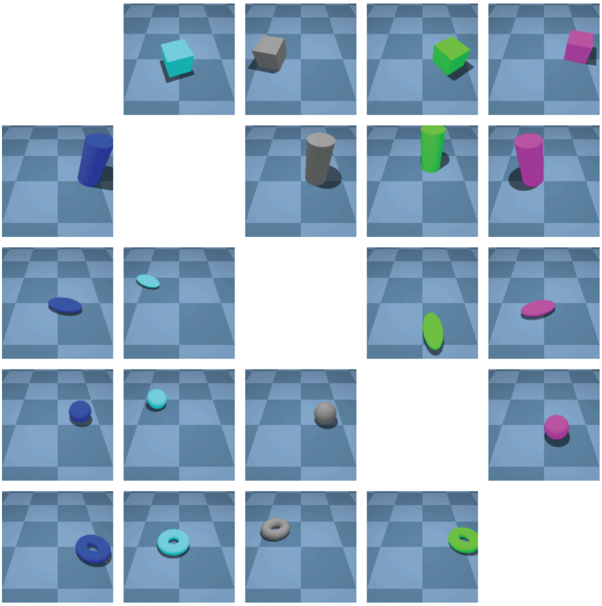

Figure 1. Examples of images from the training dataset

3.1.2 *Object naming game*

Object naming games are Lewis signaling games (Lewis, 1969; Skyrms, 2010) where the loss function $\mathcal{L}$ can be decomposed into a sum of two loss functions $\mathcal{L}_1$ and $\mathcal{L}_2$. In the object naming game used in the experiments, two agents, a sender and a receiver, learn to communicate about colored geometric objects. The sender observes an object (an RGB image) and sends a message (a sequence of discrete symbols) to the receiver; the receiver must correctly indicate both the color and the shape of the object. Formally, the game is stated as maximization of the following log likelihood:

$$\mathcal{L}(\theta, \Psi) := \mathbb{E}_{x,yc,y\,s\sim D}\mathbb{E}_{m\sim s\theta(\cdot|x)}[-\log r_{\Psi}\,(y_c, y_s|m)], \tag{1}$$

where s_θ is the policy of the sender (i.e. $s_\theta(m|x)$ is the probability of sending message m when observing image x), r_Ψ is the policy of the receiver (i.e. $r_\Psi(y_c, y_s|m)$ is the probability of taking actions y_c, y_s after receiving message m). D is the dataset, and a sample of the dataset consists of the following: x, an RGB representation of the object, and labels y_c and y_s for the color and shape of the objects. Parameters θ and Ψ are learnable parameters of the polices. For more details, see Algorithm 1 and Figure 2.

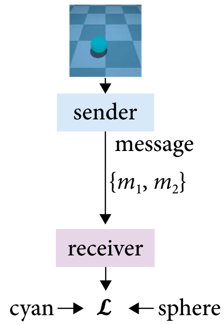

Figure 2. Object naming game

Algorithm 1 Training loop for the object naming game

Initialize sender s_θ, receiver r_Ψ, and training set D
for $x, yc, ys \in$ D do
    $m \sim s_\theta(x)$
    $\hat{y}_c, \hat{y}_s = r_\Psi(m)$
    $\mathcal{L} = $ -log_likelihood($y_c, \hat{y}_c$) – log_likelihood($y_s, \hat{y}_s$)
    optimize($\mathcal{L}(\theta, \Psi)$)

3.2 Architecture of the agents

General setup

Both the sender and the receiver are implemented as simple recurrent neural networks (Elman, 1990). The sender is equipped with a pre-trained convolutional neural network (LeCun et al., 1998) to process visual input. After observing the object, the sender generates a sequence of T discrete messages sampled from a closed vocabulary of 10 symbols. The last (softmax) layers are distinct for shape and colors, but earlier layers (including the RNNs) are shared and hence processing of shape and color will be entangled by default. In the naming games, a color sender and a shape sender produce messages, and the first symbol of color sender's message and the second symbol of shape sender's message are concatenated and then passed over to the receiver. It is in this sense that the agents are optimised simultaneously.

All experiments reported in this paper are implemented using PyTorch (Paszke et al., 2017) and EGG (Kharitonov et al., 2019). The code is publicly available.[4]

Vision module

We pre-train a simple convolutional neural network on the training subset of our datatset to predict colors and shapes. The network is composed of two layers of filters (20 and 50 filters with kernel size 5×5 and stride 1), each followed by a ReLU (rectified linear unit) activation and max pooling. The output of convolutional layers is then projected into a 25-dimensional image embedding using a fully-connected layer. During pre-training, the image embedding is passed to two linear classifiers (for color and shape) and the whole vision module is optimized with negative log likelihood as a loss function.

Sender

During naming games, the vision module is kept frozen (i.e. it is not updated during training). The sender generates its messages using a single-layer recurrent neural network (RNN) with a hidden state size of 200. The 25-dimensional image embedding for each image is projected to 200 dimensions to initialize the hidden state of the RNN. Let T be a fixed length of the message. Then, at each time-step $t < T$, the output of the RNN is used to parameterize a Gumbel-Softmax distribution (together with a temperature t that is a trainable parameter as well). A symbol is sampled from this distribution at each time-step t. After reaching T, the RNN halts and the generated symbols are concatenated to form a message, which is then passed to the receiver.

Receiver

The receiver processes a message symbol-by-symbol using a single-layer recurrent neural network with a hidden state size of 200. After processing the entire sequence, the last output is passed to a two-layer neural network classifier with two softmax outputs for color and shape.

Hyperparameters

All models are optimized using Adam (Kingma and Ba, 2014). The batch size is always 32. During the object naming game, the sender is trained with learning rate 10^{-5} and receiver with learning rate 10^{-5}.

4. The code for this paper is available at https://github.com/tomekkorbak/Interaction-history-as-a-source-of-compositionality.

3.3 Template transfer approach

The template transfer approach boils down to pre-training the receiver on two simpler guessing games: a *color naming game* and a *shape naming game.* These games are disentangled in the sense that their tasks are to correctly indicate one aspect of the object (color or shape), as formalized by the following loss functions:

$$\mathcal{L}_1(\theta_1, \Psi) := \mathbb{E}_{(x,y_c)\sim D}\, \mathbb{E}_{m\sim s\theta_1\,(\cdot|x)}[-\log r_\Psi\,(y_c|m)], \tag{2}$$

$$\mathcal{L}_2(\theta_2, \Psi) := \mathbb{E}_{(x,y_s)\sim D}\, \mathbb{E}_{m\sim s\theta_2\,(\cdot|x)}[-\log r_\Psi\,(y_s|m)], \tag{3}$$

where $r_\Psi\,(y_c|m)$ is the marginalization of $r_\Psi a\,(y_c, y_s|m)$, viz. $r_\Psi\,(y_c|m) := \Sigma_{y_s} r_\Psi(y_c, y_s|m)$. Analogously, one can define $r_\Psi\,(y_s|m) := \Sigma_{y_c} r_\Psi\,(y_c, y_s|m)$. Crucially, as far as Y_c is conditionally independent from Y_s given X, we have

$$\mathcal{L}(\theta, \Psi) = \mathcal{L}_1(\theta, \Psi) + \mathcal{L}_2(\theta, \Psi). \tag{4}$$

The loss functions $\mathcal{L}_1$ and $\mathcal{L}_2$ are optimized simultaneously (crucially with the shared parameters Ψ of the receiver) until a desired level of accuracy is met. Then, the second phase follows, in which the receiver is passed (via template transfer) to the object naming game (as described in the previous paragraph) with a new sender.

During the pre-training phase of template transfer, both sender and receiver, as well as the vision classifier, are trained with learning rate 10^{-3}. Message length $T=1$ for each sender. To prevent distribution shift with respect to message length between games, a random uniformly sampled symbol is prepended to s_1 's messages and appended to s_2's messages.[5] After, pre-training, during the object naming game, $T=2$ and the learning rate of the transferred receiver is decreased to 10^{-5}. See Figure 3 and Algorithm 2 for more details.

The communication protocol acquired in the first phase serves as a training bias in the second phase. Informally, the new sender learns to emulate messages sent by the two specialized senders of the previous phase. Our experiments reported in section Results indicate that two-phase learning is a sufficient incentive for compositionality to emerge.

5. To understand why that is necessary, note that the difference between $T=1$ and $T=2$ is the difference between the RNN having a unimodal (just visual information) previous hidden state and a multimodal (visual information and previous symbol) previous hidden state. Therefore, with $T=1$ the RNN would have no way of adapting to (not to mention utilizing) the additional linguistic modality.

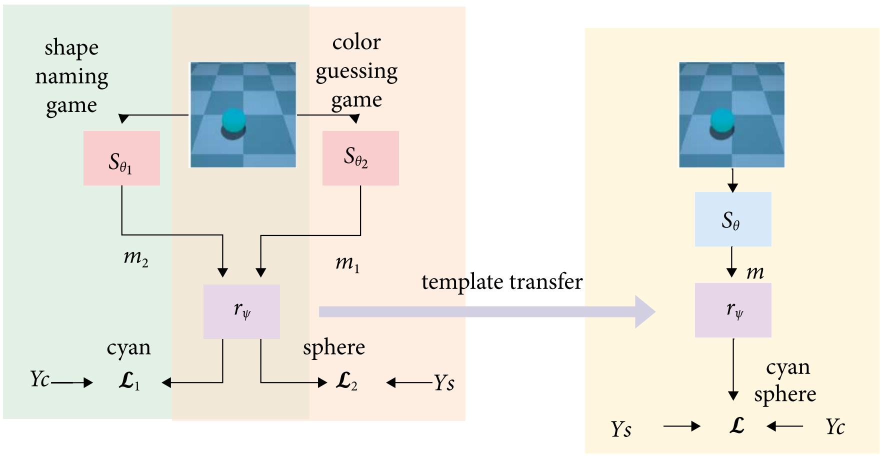

Figure 3. Template transfer consists of pre-training the receiver r_Ψ on two games with disentangled losses $\mathcal{L}_1$ and $\mathcal{L}_2$ and transferring r_Ψ to a new object naming game

Algorithm 2 Template transfer approach

Initialize senders s_{θ_1}, s_{θ_2}, s_θ, receiver r_Ψ, and training set D
for $x, y_c, y_s \in D$ **do**
$m_1 \sim s_{\theta_1}(x)$ ▷ Color naming game
$m' \sim vocabulary$
$\hat{y}_c, \hat{y}_s = r_\Psi([m_1, m'])$
$\mathcal{L}_1 =$ -log_likelihood$(y_c, \hat{y}_c)$
$m_2 \sim s_{\theta_2}(x)$ ▷ Shape naming game
$m'' \sim$ vocabulary
$\hat{x}_c, \hat{y}_s = r_\Psi([m'', m_2])$
$\mathcal{L}_2 =$ -log_likelihood$(y_s, \hat{y}_s)$
optimize$((\mathcal{L}_1(\theta_1, \Psi) + \mathcal{L}2\,(\theta_2, \Psi))$
for $x, y_c, y_s \in$ D **do**
$m \backsim s_\theta(x)$ ▷ Object naming game
$\hat{y}_c, \hat{y}_s = r_\Psi(m)$
$\mathcal{L} =$ -log_likelihood$(y_c, \hat{y}_c)$ –log_likelihood$(y_s, \hat{y}_s)$
optimize$(\mathcal{L}(\theta, \Psi))$

Our approach instantiates both template transfer and modular composition as described in section Related work. To simplify the notation, let us assume for a moment that s_θ and r_Ψ are deterministic functions $s_\theta(x) := \arg\max_m s_\theta(m|x)$ and $r_\Psi(m) := \arg\max_y \mathrm{r}_\Psi(y|m)$ (as is the case during evaluation). The fact that pre-training involves both the color naming game and the shape naming game can be seen as a modular composition of these games with a game g_a aggregating the predictions of the receiver as it communicates with each sender: g_a $(r_\Psi(s_{\theta_1}(x)), r_\Psi(s_{\theta_2}(x)))$ such that the loss for g_a is $\mathcal{L}_1 + \mathcal{L}_2$. The presented approach also instantiates template transfer from the composite game *ga* to the object naming game. In the latter game, the new sender s_θ takes advantage of the biases in the receiver r_Ψ due to playing the composite game. The significance of these interpretations in further discussed in section Discussion.

4. Experiments and results

In this section, we attempt to measure how much the template transfer approach influences the degree of compositionality of a communication protocol as compared to three baseline approaches (random agents, the same architecture without pre-training and the obverter approach). The compositionality is measured using three metrics: test accuracy, context-independence and topographical similarity, which will be described in the first section. Finally, we also try to provide an attempt at explaining how template transfer affect the biases learned by the receiver by visualizing the activations of the RNN implementing the receiver. It turns out that template transfer causes the receiver to learn disentangled representations of color and shape.

4.1 Measuring compositionality

We utilize three metrics of compositionality of a communication protocol: *zero-shot generalization accuracy, context-independence* and *topographical similarity.* High zero-shot generalization indicates that the agents correctly map the implicit compositional structure of inputs to explicate one of the outputs. The other two metrics focus directly on the transmitted messages, comparing them to the ground truth, fully disentangled (color, shape) representation.

During evaluation we use the deterministic sender given by $s(x) := \arg\max_m s_\theta(m|x)$, where x is an object.

Test set accuracy

We quantify zero-shot generalization by measuring the accuracy of the agents on a test set obtained by a compositional split of the dataset: the test set only containing pairs of shapes and colors not present in the training set, but each color and shape individually is present in the training set. Test set accuracy therefore measures the ability to generalize to unseen combinations of seen elements.

Context-independence

Context-independence was introduced by Bogin et al. (2018) as a measure of alignment between the symbols in an agent's messages and the concepts transmitted. We denote by V the set of symbols used to compose messages and by K the set of concepts, which in our case is the union of available colors and shapes. Given sender s, and assuming a uniform distribution of objects, we define $p(v|k)$ as the probability that symbol $v \in V$ appears when the sender observes an object with property $k \in K$. We define $p(k|v)$ in the same manner.[6] Further, let $v^k :=$ arg max$_v$ $p(k|v)$. The context-independence metric is defined as $\mathbb{E}(p(v^k|k) \cdot p(k|v^k))$; the expectation is taken with respect to the uniform distribution on K.

Intuitively, context-independence measures the consistency associating symbols with shapes irrespective of color (and vice versa). It is sometimes considered restrictive, as it effectively punishes for using synonyms (Lowe et al., 2019).

Topographical similarity

Finally, we introduce topographical similarity (Brighton and Kirby, 2006; Lazaridou et al., 2018), also known as *representational similarity* (Kriegeskorte, 2008; Bouchacourt and Baroni, 2018), a measure of structural similarity between messages and disentangled target labels y_c, y_s. To define topographical similarity more formally, let us denote the random variable $L_t := L((y_c^1, y_s^1), (y_c^2, y_s^2))$, where L is the Levenshtein (1966) distance and y_c^1, y_s^1 and y_c^2. y_s^2 are ground truth labels for independently objects $x^{(1)}$, $x^{(2)}$ with the subscripts denoting their shapes and colors. Note that in our case $L_t \in \{0, 1, 2\}$. Let $L_m := L(s(x^1), s(x^2))$ be the distance between messages sent by the sender after observing x^1 and x^2. Topographical similarity is the the Spearman p correlation of L_t and L_m.

Topographical similarity is theoretically principled because an analogous metric is used in computational neuroscience to measure, for instance, the structural similarity between a stimulus and neural activity evoked by the stimulus (Kriegeskorte, 2008). Moreover, being a second-order relation between the messages and ground truth labels, topographical similarity mirrors the idea of Deacon (1998) about symbolic reference being a second-order relation between indexical signs.

---

6. While Bogin et al. (2018) estimate these two conditionals based on IBM Model 1 (Brown et al., 1993), we simply compute occurrence frequencies of k given v and v given k in all recorded messages.

4.2 Baselines

To establish sensible lower bounds on all three described metrics, we measure the performance of three baseline models.

Random baseline

Random baseline is simply the performance of untrained agents.

Same architecture

The most direct comparison of the effect of template transfer is simply not applying template transfer, i.e. not pretraining the sender on color naming game and shape naming game and only training the agents on the object naming game.

Obverter baseline

In the obverter algorithm, two agents exchange the roles of the sender and the receiver. If an agent is the receiver, it behaves as in the object naming game. If an agent is the sender, it sends message that would have produced the most accurate prediction of color and shape, if it had received such a message as a receiver (i.e. instead of the greedy decoding used in the original implementation of Batali (1998), we simply choose the message maximizing accuracy). Accuracy is evaluated against the predictions of the visual classifier. The receiver is trained with learning rate 10^{-5}. For details, consult Algorithm 3.

Algorithm 3 Obverter

Initialize agents a_1, a_2, visual module v, training set D
Initialize the set M of all possible messages m
for $x, y_c, y_s \in D$ do
    $s_\theta, r_\Psi \sim \{a_1, a_2\}$ ▷ Randomly assigning the roles of sender and receiver
    $m = \arg\min_{m \in M}$ evaluate_message(s_θ, m)
    $\hat{y}_c, \hat{y}_s = r_\Psi(m)$
    $\mathcal{L} = -$ log_likelihood($y_c, \hat{y}_c$) $-$ log_likelihood ($y_s, \hat{y}_s$)
    optimize($\mathcal{L}(\Psi)$)
procedure evaluate_message (model, m)
    $y_c, y_s = v(x)$ ▷ Using visual classifier predictions as a proxy for labels
    $\hat{y}_c, \hat{y}_s =$ model (m)
    $\mathcal{L}' = -$ log likelihood($y_c, \hat{y}_c$) $-$ log_likelihood($y_s, \hat{y}_s$)
    return $\mathcal{L}'$

4.3 Results

We compared our approach with several baselines (random, the same architecture without pre-training games, and our implementation of the obverter approach) on games with five shapes and five colors. Topographical similarity and context-independence were computed on the full dataset (train and test); objects in the dataset were sampled uniformly. The results are presented in Table 1. Template transfer clearly leads to highly compositional communication protocols. While all methods struggled to generalize to unseen objects, template transfer was the most successful. The relatively strong performance of the baseline model in terms of average accuracy is perhaps surprising: it correctly guesses the color 47% of the time (on average) and correctly guesses the shape 47% of the time (on average), but guesses correctly both only 2% of the time. This is because of an anti-correlation between shape and color prediction, i.e. for various objects, the agents specialized in classifying correctly its shape or color. Thus, the "both" column is more meaningful for comparing accuracies: there is a gap between two baselines and Obverter on the one hand, and between Obverter and template transfer on the other.

Table 1. The effect of template transfer on compositionality. The metrics are train and test set accuracies (the rate of correctly predicted both y_c and y_s); average over the individual accuracies for y_c and y_s; and context-independence (CI) and topographical similarity (Topo). The models are random baseline (untrained agents); baseline architecture (without template transfer); template transfer (TT); and obverter algorithm. All reported metrics are averaged over ten random seeds and standard deviations are reported in brackets

	Accuracy				
Model	**Train (both)**	**Test (both)**	**Test (avg)**	**CI**	**Topo**
Random	0.04	0.04	0.2	0.04 (± 0.01)	0.13 (± 0.03)
Baseline	0.99 (± 0.01)	0.02 (± 0.05)	0.47 (± 0.09)	0.08 (± 0.01)	0.30 (± 0.05)
Obverter	0.99 (± 0)	0.24 (± 0.23)	0.51 (± 0.19)	0.12 (± 0.02)	0.55 (± 0.13)
Template transfer	1(± 0)	0.48 (± 0.10)	0.74 (± 0.06)	0.18 (± 0.01)	0.85 (± 0.03)

For examples of communication protocols representative of the experiments conducted, see Table 2 and Figure 4. We can observe that with template transfer each shape (color) is systematically associated with the same symbol across multiple colors (shapes). In contrast, in the baseline model the symbol associated with each color (shape) changes across shapes (colors).

4.4 Visualizing the activations of the receiver

Recall that the receiver r_Ψ consists of an RNN that reads the message symbol-by-symbol and a two-layer neural network classifier with two softmax heads: one for color and one for shape. The last hidden state of the RNN serves an input to the two-layer neural network classifier. To get a better sense of how the receiver understands the messages it receives, we visualized the hidden states h_m corresponding to each message m sent by the sender after receiving each object x, sampling one object for each color-shape pair. Then we applied principal component analysis, computed a projection $\text{proj}_{p1,\,p2}\, h_m$ of each hidden state onto two principal components p_1 and p_2.[7] The scatter plots visualizing the RNN hidden states for the baseline architecture and template transfer are shown on Figure 5.

While without template transfer there is no clear structure in the space, the RNN of the receiver trained with template transfer exhibits clear structure: color and shape are linearly separable and spanned by the two principal components of the representation space. One can observe that representations learned by the receiver are disentangled in the sense that the features within the representation correspond to the underlying causes of the observed data, with separate features corresponding to different causes (Goodfellow et al., 2016). The causes in our case are color and shape. Since disentanglement can be seen as a representational correlate of compositionality, it provides further evidence that the semantics agents use to produce and comprehend messages is indeed compositional (i.e. there is semantic compositionality in addition to syntactic compositionality).

Table 2. Two example communication protocols, one that emerged via the baseline architecture (2a), and one via template transfer (2a). Gray cells indicate objects not seen during training. In (2b), symbols exhibit clear association with colors and shapes, e.g. symbol 8 is consistently associated with the color magenta (when on first position) and boxes (when on second position)

a. A non-compositional communication protocol (topographical similarity 0.25)

	Box	**Sphere**	**Cylinder**	**Torus**	**Ellipsoid**
blue	10	45	10	45	50
cyan	90	40	30	40	70
gray	35	65	32	65	53
green	00	76	30	60	76
magenta	15	55	12	15	52

7. It is common in the literature to use PCA as a method for visualizing hidden states of RNNs, see also Yamashita and Tani (2008).

b. A highly compositional communication protocol (topographical similarity 0.85)

	Box	Sphere	Cylinder	Torus	Ellipsoid
blue	18	19	15	16	14
cyan	48	49	45	46	44
gray	68	69	65	66	69
green	98	99	95	96	94
magenta	88	89	85	88	84

5. Discussion

In this section, we discuss the implications of our approach focusing on the following points: (i) potential sources of compositionality, (ii) how compositionality can be understood game-theoretically, (iii) how the results corroborate Deacon's (1998) account of reference, (iv) that compositionality may be less cognitively demanding that previously thought, and (v) that the presented approach is developmentally plausible to an extent.

Sources of compositionality

In Lewis signaling game pressure for the sender to produce compositional messages may come from the environment structure, from its own innate biases or from the dispositions of the receiver. We use a relatively simple environment and a general agent architecture, which do not put much pressure on the sender. The pressure for compositionality comes from the receiver. When the receiver is highly flexible and is able to pick the sender's meaning perfectly it does not introduce any pressure at all. In our work it is the opposite – the receiver is highly specific due to its previous history of interactions in a changing context. Two aspects are worth noting: i) the receiver is prepared to understand compositional messages even before the first compositional message is produced, ii) the sender learns to produce compositional messages describing objects after the initial reason for disentangling object properties disappears.

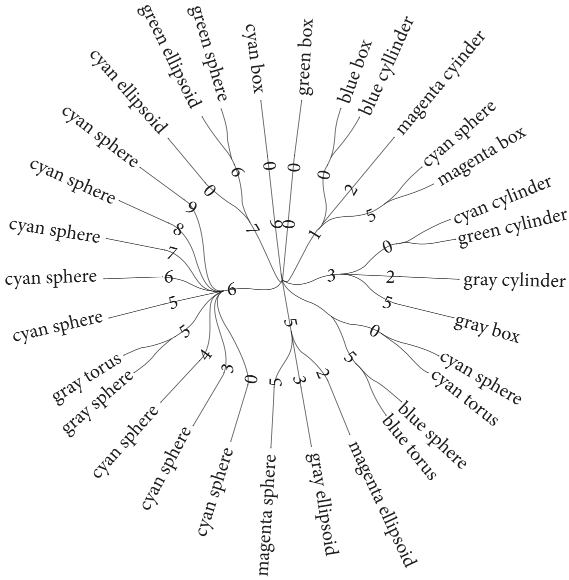

a. A non-compositional communication protocol (topographical similarity 0.25)

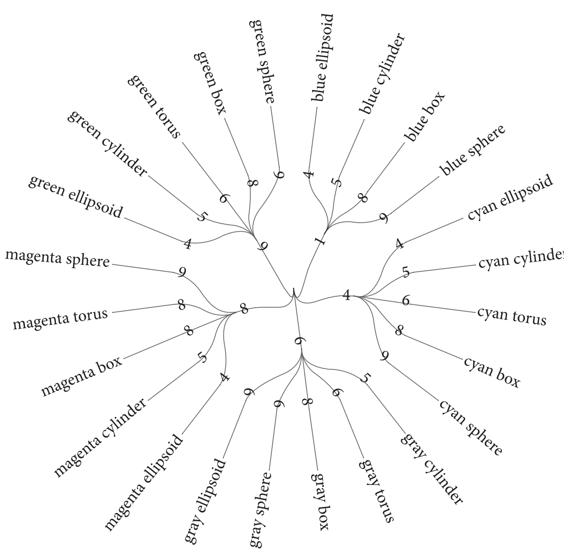

b. A highly compositional communication protocol (topographical similarity 0.85)

Figure 4. Communication protocols in the object naming game admit an information-theoretic interpretation as prefix code, which can be visualized as a tree. Here we visualize the trees corresponding to a non-compositional protocol and a high-compositional protocol. Note that compositionality – which can be seen as a kind of symmetry in the protocol – is depicted by radial symmetry of the corresponding tree

We demonstrate this process for the object naming game with a disembodied receiver, but the same set of pressures might be recreated under different setups. Let us consider a model in which the receiver is presented with a set of objects and has to pick the correct one. There, we could manipulate the distribution of distractor objects in such a way as to effectively reward only distinguishing shapes or colors. This would lead to a disentangled loss required for the template transfer.

We expect that in the real world multiple pressures affect sender and receiver simultaneously. History of receiver's interactions is just one source of such pressures. It is unique because it is able to explain compositionality as a consequence of past events, regardless of whether it is functional under the present circumstances.

Evolutionary game-theoretic interpretation

The presented approach instantiates both template transfer (from the pre-training games to the object naming game) and modular composition (of the color naming game and the shape naming game), concepts developed by Barrett and Skyrms (2017; 2019) and described in section

Related work. The subsequent discussion focuses on template transfer as a mechanism for reusing skills across contexts and scaffolding compositional communication protocols with simpler protocols. The exploitation of modular composition, however, also offers a theoretical insight. There is an interesting analogy between symbolic composition, an operation over symbols yielding composite symbols, and modular composition, an operation over games yielding composite games. The assumption central to our approach is that a game (such as the object naming game) can be reformulated as a modular composition of two simpler games: $g_a(r_\Psi(s_{\theta_1}(x)), r_\Psi(s_{\theta_2}(x)))$ with a game g_a being a function aggregating the predictions of color and shapes and sender and receiver assumed to be deterministic for notational convenience. Under this formulation, we can have specialized senders s_{θ_1} and s_{θ_2} for the pre-training games. Therefore, decomposing a game – and enabling the agent to specialize in sub-games – is sufficient for compositionality to emerge. One could then conjecture that compositional communication is itself a composition of distinct communication skills and as such it follows a more basic kind of compositionality – composing simple skills to give rise to complex behavior. That conjecture fits well with the account of skill reuse in nature formalized as generalized signaling games and would place our model of compositional communication among evolutionary game-theoretic explanations of social phenomena such as inference (Barrett and Skyrms, 2017), knowledge sharing (Barrett et al., 2019) and functional specialization (Barrett et al., 2018).

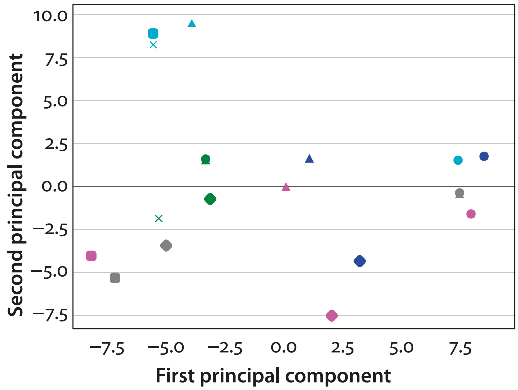

a. No template transfer

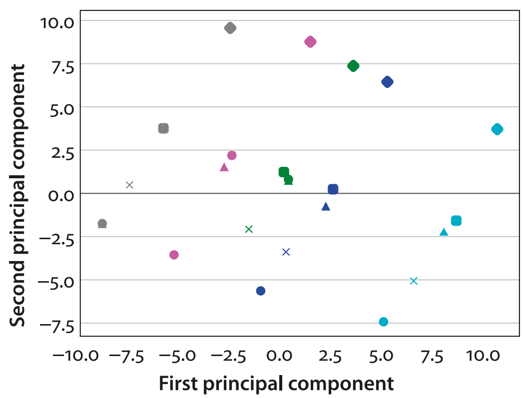

b. Template transfer

Figure 5. Receiver RNN's hidden states corresponding for each object type plotted on a 2d plane. Scatter point indicate indicate inputs to the sender, their colors and shapes indicate the color (blue, cyan, gray, green, magenta) and shapes (box, sphere, cylinder, torus ellipsoids) of corresponding inputs

Semiotic interpretation

Peirce (1998) famously proposed a hierarchy of forms of signification:

1. Iconic signs refer to their objects by virtue of physical similarity between a sign and an object as perceived by an agent,
2. Indexical signs refer to their objects by virtue of causal, spatial or temporal association between a sign and an object as recognized by an agent,

3. Symbolic signs refer to their objects by virtue of a social convention or tacit agreement familiar to an agent.

Peirce's account of precedence and dependence of different forms of reference is influential both in evolutionary research on the origins of language as well as in language development research. It is frequently assumed as a target evolutionary pathway in computational models of the evolution of language (Cangelosi, 2001; Grouchy et al., 2016).

Deacon (1998) developed a cognitive anthropological interpretation of Peirce's semiotics and argued that the linear order over three kinds of signs is to be interpreted both in terms of ontogenetic and phylogenetic precedence as well as evolutionary and developmental functional dependence. Regarding precedence, the hierarchy reflects an ascending order of cognitive competence required to interpret respective signs. Iconic reference requires modest cognitive capacities to be recognized (perception fine-grained enough to recognize similarity, but without the requirement for memory) while indexical reference requires a form of associative learning. Finally, symbolic reference requires reasoning according to rules defined by a whole system of symbols (Peirce, 1998).[8] Empirically, sensitivity to iconic reference can be found arbitrarily early in phylogeny and most animal communication systems are indexical. Symbolic reference is usually assumed to be unique to human languages (Deacon, 1998).

There is, however, another view on Peirce's hierarchy according to which the order should be taken not as (evolutionary, developmental or cognitive) precedence, but as a part-of relationship. According to Deacon, "reference is hierarchical in nature; more complex forms of reference are built up from simpler forms" (Deacon, 1998, p. 73). This is because the competence to interpret symbolically assumes competence to interpret indexically (and by consequence, iconically). In Peirce's own terms, higher-order forms of reference can be decomposed into lower order forms in the sense that a lower order form of reference (e.g. an icon) serves as an interpretant to a higher order form (e.g. an index).

Inspired by Peirce and Deacon, our approach solves the problem of developing compositional communication protocol from raw pixel input by decomposing the problem into several simpler problems. These simpler problems are:

1. Learning a visual classifier,
2. Learning non-compositional communication protocols in simple games, and

8. This is because symbolic reference between a symbol *S* and an object *O* is determined by a relationship *S* has with other symbols *S'*, *S''*…, not just the relationship between *S* and *O*. Relationships between symbols may involve rules for composing them, e.g. certain co-occurrences are allowable and others forbidden.

3. Learning a compositional communication protocol.

These problems map into the hierarchy different forms of reference: iconic reference, indexical reference and complex indexical reference[9] (Peirce, 1998; Deacon, 1998). Thus, template transfer implements the Peircean conjecture that compositional communication is preceded (both evolutionarily and developmentally) by progressively augmented iconic and indexical communication protocols. It also illustrates how the idea of simpler forms of reference used as a scaffolding for complex forms of reference can be formalized in terms Lewis signaling games by appealing to modular composition and template transfer. More specifically, both the color naming game and the shape naming game considered separately instantiate simple indexical communication between s_{θ_1} and r_Ψ and between s_{θ_2} and r_Ψ.

Additionally, the pre-training game (composed of the color naming game and the shape naming game) constrains the receiver to interpret the messages of *both* s_{θ_1} and s_{θ_2} compositionally. It is this inductive bias – the receiver playing the role of a compositional interpretant (in Peirce's sense) – that further constrains the new sender s_θ to communicate compositionally.

Cognitive interpretation

Some of the existing methods of inducing compositionality (e.g. the obverter approach) focus on imposing strong inductive biases on the architecture of the agents (recall section Related work). For instance, the obverter approach is based on the assumption that an agent can use its own responses to messages to predict other agent's responses and thus can iteratively compose its messages to maximize the probability of desired response (according to the self-model). Therefore, it makes strong assumptions about the agents and task: to be able to use themselves as models of others, the agents must share an identical architecture and the task must be symmetric (the agent must be able to exchange their roles). This excludes games with functional specialization of agents. Template transfer is a model-free technique that makes one assumption (similar to the one present in Nowak and Krakauer (1999)): that the loss function can be decomposed into two disentangled loss functions (as in the case of decomposing $\mathcal{L}$ into $\mathcal{L}_1$ and $\mathcal{L}_2$ in (2)–(3). (Note that there is no need for the input to be disentangled.) The fact that template transfer can outperform the obverter approach on all compositionality met-

9. We deliberately refrain from using the term "symbolic reference" because the communication protocols learned by the agents in presented experiments are not non-controversially symbolic in the sense of Deacon as they lack rich symbol-to-symbol relationships.

rics lends support to the claim that the cognitive requirements for developing a compositional communication protocol are quite modest.

One may argue that template transfer offloads some of the cognitive complexity of learning a compositional communication protocol to the interaction history, supporting a distributed view of language as an activity happening in a social world that evolves outside of individual speakers (Cowley, 2011). According to the distributed view of language, a speaker might be constrained by multiple interactions antecedent to the speaker coming to being. The usefulness of this view is more evident when thinking about the compositional communication protocol that the receiver learns in the pre-training game as instantiating a replicable constraint in the sense of (Rączaszek-Leonardi, 2012). Assuming this picture, language is a system of physical structures that act as constraints, selected due to having a history of harnessing dynamics in a useful way and transmitted between settings. Importantly, the emergence and transmission usually happen on a slower timescale than the actual constraining. In the conducted experiments, the compositional communication protocol was learned in the pre-training phase as a useful way of harnessing the communication dynamics. Due to its usefulness, it persisted in receiver's weights, which allowed it to replicate to the object naming game, constraining a new sender via receiver's expectations. In effect, the new sender s_θ took advantage of the solution to the coordination problem developed jointly by s_{θ_1}, s_{θ_2}, and inherited it implicitly, never interacting with s_{θ_1} and s_{θ_2}. This illustrates how the communication constraints emerge in a distributed way in a structured social environment and depend on each other (in a sense that pre-existing constraints unleash novel forms of communication) (Raczaszek-Leonardi et al., 2018). It is in this sense that the problem of learning to communicate compositionally can be solved much more easily by agents embedded in a rich, social world.

Developmental interpretation

While solving the problem of developing compositional communication protocol from raw pixel input and learning compositional communication from scratch in an end-to-end manner (Lazaridou et al., 2018; Choi et al., 2018) is of theoretical interest, it significantly differs from how human children learn compositional aspects of language. Children learn communicative functions of utterances in a rich and highly structured environment (child-directed speech exhibits repetitive patterns and is augmented with pointing, gazing or other means of attention shifting) and through simple language games that lack many features of adult language (Stern, 1974; Bruner, 1983; Nomikou et al., 2017; Raczaszek-Leonardi et al., 2018). The template transfer approach is developmentally inspired as it acknowledges both the piecemeal (children learn words holistically before learning complex syntactical constructions) and the socially embedded (the role of child-directed speech) character of language development.

6. Conclusions

The goal of the paper was to present a novel approach to developing emergent compositional communication based on the idea of template transfer (Barrett and Skyrms, 2017) implemented by sharing agents across games. Template transfer was used to model a variety of semiotic, social and cognitive phenomena (Barrett et al., 2019, 2018) and can probably be extended to new, more challenging problems in multi-agent systems research.

Limitations

Our model is limited by the simplicity of the task and the static nature of the environment. The communication channel is constrained by predefined vocabulary size (10) and message length (2), and further by partitioning the channel in the pre-training game into single-symbol subchannels for sender s_{θ_1} and s_{θ_2}. Messages from the two senders are presented to the receiver at fixed positions in a sequence (s_{θ_1}'s message at first position, s_{θ_2}'s message at second position). This potentially simplifies the task and forces a fixed order of symbols in the evolved compositional language. Moreover, there are only two effective degrees of freedom in the world (color and shape), agents are assigned with specific roles and they do not control which object they are being presented with. Future work might focus on extending the template transfer approach to more realistic, interactive 3d environments with messages of arbitrary length, non-trivial compositionality (Steinert-Threlkeld, 2020; Korbak et al., 2020) could emerge. A richer structure of the environment and the task could also lead to the emergence of symbolic reference (Deacon, 1998) with the meanings of messages being deeply interconnected.

Funding

Tomasz Korbak was funded by Ministry of Science and Higher Education (Poland) grant DI2015010945. Joanna Raczaszek-Leonardi and Julian Zubek were funded by National Science Centre (Poland) grant OPUS 2018/29/B/HS1/00884. Piotr Miłoś was funded by National Science Centre (Poland) grant UMO-2017/26/E/ST6/00622. The computational experiments presented in this paper were run on the Prometheus cluster as part of the PLGRID infrastructure.

Acknowledgements

The authors are grateful to Jakub Szymanik, Travis LaCroix, Jeffrey Barrett and the participants of Human Interactivity and Language Lab (Faculty of Psychology, University of Warsaw) seminars, the 8th Peripatetic Conference on Modeling Cognitive Systems (Kiry, Poland, October 2019), PL in ML conference (Warsaw, Poland, November 2019) and NeurIPS 2019 workshop Emergent communication: Towards Natural Language (Vancouver, Canada, December 2019) for their valuable comments on early results leading to this paper. Finally, the authors are grateful to two anonymous reviewers for their helpful feedback.

References

Barrett, J., Skyrms, B., and Cochran, C. (2018). Hierarchical Models for the Evolution of Compositional Language. *Institute for Mathematical Behavioral Sciences Technical Report MBS 18-03*.

Barrett, J.A. and Skyrms, B. (2017). Self-assembling Games. *The British Journal for the Philosophy of Science*, *68*(2):329–353. https://doi.org/10.1093/bjps/axv043

Barrett, J.A., Skyrms, B., and Mohseni, A. (2019). Self-Assembling Networks. *The British Journal for the Philosophy of Science*, *70*(1):301–325. https://doi.org/10.1093/bjps/axx039

Batali, J. (1998). Computational simulations of the emergence of grammar. *Approach to the Evolution of Language*, pages 405–426.

Bogin, B., Geva, M., and Berant, J. (2018). Emergence of Communication in an Interactive World with Consistent Speakers. *arXiv:1809.00549 [cs]*. arXiv: 1809.00549.

Bouchacourt, D. and Baroni, M. (2018). How agents see things: On visual representations in an emergent language game. *arXiv:1808.10696 [cs]*. arXiv: 1808.10696. https://doi.org/10.18653/v1/D18-1119

Brighton, H. (2002). Compositional Syntax From Cultural Transmission. *Artificial Life*, *8*(1):25–54. https://doi.org/10.1162/106454602753694756

Brighton, H. and Kirby, S. (2006). Understanding Linguistic Evolution by Visualizing the Emergence of Topographic Mappings. *Artificial Life*, *12*(2):229–242. https://doi.org/10.1162/artl.2006.12.2.229

Brown, P.F., Della-Pietra, S.A., Della-Pietra, V.J., and Mercer, R.L. (1993). The mathematics of statistical machine translation. *Computational Linguistics*, *19*(2):263–313.

Bruner, J.S. (1981). Intention in the structure of action and interaction. *Advances in Infancy Research*, *1*:41–56.

Bruner, J.S. (1983). *Child's talk: learning to use language*. W.W. Norton, New York, 1st ed edition.

Cangelosi, A. (2001). Evolution of communication and language using signals, symbols, and words. *IEEE Transactions on Evolutionary Computation*, *5*(2):93–101. https://doi.org/10.1109/4235.918429

Chaabouni, R., Kharitonov, E., Bouchacourt, D., Dupoux, E., and Baroni, M. (2020). Compositionality and Generalization In Emergent Languages. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4427–4442, Online. Association for Computational Linguistics. https://doi.org/10.18653/v1/2020.acl-main.407

Choi, E., Lazaridou, A., and de Freitas, N. (2018). Compositional Obverter Communication Learning From Raw Visual Input. *ICLR 2018*. arXiv: 1804.02341.

Cogswell, M., Lu, J., Lee, S., Parikh, D., and Batra, D. (2019). Emergence of Compositional Language with Deep Generational Transmission. *ArXiv*, abs/1904.09067.

Cowley, S. J., editor (2011). *Distributed language*. Number v. 34 in Benjamins current topics. John Benjamins Pub. Co, Amsterdam; Philadelphia. OCLC: ocn741355729. https://doi.org/10.1075/bct.34

Das, A., Kottur, S., Moura, J. M. F., Lee, S., and Batra, D. (2017). Learning Cooperative Visual Dialog Agents with Deep Reinforcement Learning. *2017 IEEE International Conference on Computer Vision (ICCV), Venice, 2017*. arXiv: 1703.06585. https://doi.org/10.1109/ICCV.2017.321

De Beule, J. and Bergen, B. K. (2006). On the emergence of compositionality. In *The Evolution of Language*, pages 35–42, Rome, Italy. World scientific. https://doi.org/10.1142/9789812774262_0005

Deacon, T. W. (1998). *The symbolic species: the co-evolution of language and the brain*. Norton, New York, NY, norton paperback edition. OCLC: 254499872.

Elman, J. L. (1990). Finding Structure in Time. *Cognitive Science*, *14*(2):179–211. https://doi.org/10.1207/s15516709cog1402_1

Foerster, J. N., Assael, Y. M., de Freitas, N., and Whiteson, S. (2016). Learning to Communicate with Deep Multi-Agent Reinforcement Learning. *NIPS'16 Proceedings of the 30th International Conference on Neural Information Processing Systems*. arXiv: 1605.06676.

Gong, T., Shuai, L., and Zhang, M. (2014). Modelling language evolution: Examples and predictions. *Physics of Life Reviews*, *11*(2):280–302. https://doi.org/10.1016/j.plrev.2013.11.009

Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep learning*. Adaptive computation and machine learning. The MIT Press, Cambridge, Massachusetts.

Grouchy, P., D'Eleuterio, G. M. T., Christiansen, M. H., and Lipson, H. (2016). On The Evolutionary Origin of Symbolic Communication. *Scientific Reports*, 6:34615. https://doi.org/10.1038/srep34615

Jaques, N., Lazaridou, A., Hughes, E., Gulcehre, C., Ortega, P. A., Strouse, D. J., Leibo, J. Z., and de Freitas, N. (2018). Social Influence as Intrinsic Motivation for Multi-Agent Deep Reinforcement Learning. *arXiv:1810.08647 [cs, stat]*. arXiv: 1810.08647.

Kharitonov, E. and Baroni, M. (2020). Emergent Language Generalization and Acquisition Speed are not tied to Compositionality. *arXiv:2004.03420 [cs]*. arXiv: 2004.03420. https://doi.org/10.18653/v1/2020.blackboxnlp-1.2

Kharitonov, E., Chaabouni, R., Bouchacourt, D., and Baroni, M. (2019). EGG: a toolkit for research on Emergence of lanGuage in Games. *arXiv:1907.00852 [cs]*. arXiv: 1907.00852. https://doi.org/10.18653/v1/D19-3010

Kingma, D. P. and Ba, J. (2014). Adam: A Method for Stochastic Optimization. *arXiv:1412.6980 [cs]*. arXiv: 1412.6980.

Kirby, S. (2001). Spontaneous evolution of linguistic structure-an iterated learning model of the emergence of regularity and irregularity. *IEEE Transactions on Evolutionary Computation*, *5*(2):102–110. https://doi.org/10.1109/4235.918430

Kirby, S., Cornish, H., and Smith, K. (2008). Cumulative cultural evolution in the laboratory: An experimental approach to the origins of structure in human language. *Proceedings of the National Academy of Sciences*, *105*(31):10681–10686. https://doi.org/10.1073/pnas.0707835105

Korbak, T., Zubek, J., and Raczaszek-Leonardi, J. (2020). Measuring non-trivial compositionality in emergent communication. *arXiv preprint arXiv:2010.15058*.

Kottur, S., Moura, J. M. F., Lee, S., and Batra, D. (2017). Natural Language Does Not Emerge 'Naturally' in Multi-Agent Dialog. *arXiv:1706.08502 [cs]*. arXiv: 1706.08502. https://doi.org/10.18653/v1/D17-1321

Kriegeskorte, N. (2008). Representational similarity analysis – connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*. https://doi.org/10.3389/neuro.06.004.2008

Kuciński, L., Korbak, T., Kolodziej, P., and Miloś, P. (2021). Catalytic role of noise and necessity of inductive biases in the emergence of compositional communication. In *Advances in Neural Information Processing Systems*.

Lake, B. M., Ullman, T. D., Tenenbaum, J. B., and Gershman, S. J. (2016). Building Machines That Learn and Think Like People. *arXiv:1604.00289 [cs, stat]*. arXiv: 1604.00289.

Lazaridou, A., Hermann, K. M., Tuyls, K., and Clark, S. (2018). Emergence of Linguistic Communication from Referential Games with Symbolic and Pixel Input. *arXiv:1804.03984 [cs]*. arXiv: 1804.03984.

Lazaridou, A., Peysakhovich, A., and Baroni, M. (2016). Multi-Agent Cooperation and the Emergence of (Natural) Language. *arXiv:1612.07182 [cs]*. arXiv: 1612.07182.

LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. (1998). Gradient-based learning applied to document recognition. In *PROCEEDINGS OF THE IEEE*, pages 2278–2324. https://doi.org/10.1109/5.726791

Levenshtein, V. I. (1966). Binary Codes Capable of Correcting Deletions, Insertions and Reversals. *Soviet Physics Doklady*, *10*:707.

Lewis, D. K. (1969). *Convention: a philosophical study*. Blackwell, Oxford, nachdr. edition. OCLC: 837747718.

Li, F. and Bowling, M. (2019). Ease-of-Teaching and Language Structure from Emergent Communication. *arXiv:1906.02403 [cs]*. arXiv: 1906.02403.

Lowe, R., Foerster, J., Boureau, Y.-L., Pineau, J., and Dauphin, Y. (2019). On the Pitfalls of Measuring Emergent Communication. *arXiv:1903.05168 [cs, stat]*. arXiv: 1903.05168.

Nomikou, I., Koke, M., and Rohlfing, K. J. (2017). Verbs in Mothers' Input to Six-Month-Olds: Synchrony between Presentation, Meaning, and Actions Is Related to Later Verb Acquisition. *Brain Sciences*, *7*(12):52. https://doi.org/10.3390/brainsci7050052

Nowak, M. A. and Krakauer, D. C. (1999). The evolution of language. *Proceedings of the National Academy of Sciences*, *96*(14):8028–8033. Publisher: National Academy of Sciences Section: Biological Sciences. https://doi.org/10.1073/pnas.96.14.8028

Oliphant, M. and Batali, J. (1997). Learning and the Emergence of Coordinated Communication. *Center for Research on Language Newsletter*, *11*.

Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A. (2017). Automatic differentiation in PyTorch.

Peirce, C. S. (1998). *Collected papers of Charles Sanders Peirce*. Thoemmes Press, Bristol, England. OCLC: ocm39692049.

Raczaszek-Leonardi, J. (2012). Language as a System of Replicable Constraints. In *LAWS, LANGUAGE and LIFE*, volume 7, pages 295–333. Springer Netherlands, Dordrecht. https://doi.org/10.1007/978-94-007-5161-3_19

Raczaszek-Leonardi, J., Nomikou, I., Rohlfing, K. J., and Deacon, T. W. (2018). Language Development From an Ecological Perspective: Ecologically Valid Ways to Abstract Symbols. *Ecological Psychology*, *30*(1):39–73. https://doi.org/10.1080/10407413.2017.1410387

Rainey, H. J., Zuberbühler, K., and Slater, P. J. B. (2004). Hornbills can distinguish between primate alarm calls. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, *271*(1540):755–759. https://doi.org/10.1098/rspb.2003.2619
Rumelhart, D. E., Hinton, G. E., and Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, *323*(6088):533–536. https://doi.org/10.1038/323533a0
Ryan, M. and Rand, S. (1993). Sexual selection and signal evolution: the ghost of biases past. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, *340*(1292):187–195. https://doi.org/10.1098/rstb.1993.0057
Skyrms, B. (2010). *Signals: evolution, learning, & information*. Oxford University Press, Oxford; New York. OCLC: ocn477256653. https://doi.org/10.1093/acprof:oso/9780199580828.001.0001
Steinert-Threlkeld, S. (2020). Towards the Emergence of Non-trivial Compositionaliy. *Philosophy of Science*. https://doi.org/10.1086/710628
Stern, D. N. (1974). The Goal and Structure of Mother-Infant Play. *Journal of the American Academy of Child Psychiatry*, *13*(3):402–421. https://doi.org/10.1016/S0002-7138(09)61348-0
Tomasello, M., Carpenter, M., Call, J., Behne, T., and Moll, H. (2005). Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and Brain Sciences*, *28*(5):675–691. https://doi.org/10.1017/S0140525X05000129
Wittgenstein, L. (1953). *Philosophical Investigations*. Basil Blackwell, Oxford.
Yamashita, Y. and Tani, J. (2008). Emergence of Functional Hierarchy in a Multiple Timescale Neural Network Model: A Humanoid Robot Experiment. *PLoS Computational Biology*, *4*(11):e1000220. https://doi.org/10.1371/journal.pcbi.1000220

Address for correspondence

Tomasz Korbak
Department of Informatics
University of Sussex
Chichester 1
Falmer, Brighton BN1 9RH
United Kingdom
tomasz.korbak@gmail.com

Biographical notes

Tomasz Korbak is a PhD student at the Department of Informatics, University of Sussex working on deep reinforcement learning, generative models and computational neuroscience with Dr. Chris Buckley and Prof. Anil Seth. Previously, he was a researcher at the Institute of Philosophy and Sociology, Polish Academy of Sciences working with Prof. Marcin Miłkowski and a research assistant at the Human Interactivity and Language Lab, Faculty of Psychology, University of Warsaw working with Prof. Joanna Rączaszek-Leonardi. He is interested in biologically-inspired and probabilistic approaches to artificial intelligence and recently focuses on emergent communication, compositionality, control as inference and controllable generative modeling.

Julian Zubek is an assistant professor at the Department of Psychology, University of Warsaw and a researcher at the Human Interactivity and Language Lab. He received PhD in computer science from the Institute of Computer Science, Polish Academy of Science. His research interests include cognitive processes modeling, dynamical systems, embodied and enactive theories of human cognition. He applies machine learning methods and statistical models in the context of experimental psychology.
j.zubek@uw.edu.pl

Łukasz Kuciński is an Assistant Professor at the Polish Academy of Sciences. His research focuses on machine learning and he co-leads a reinforcement learning lab. He contributes to the growth of the Polish machine learning community, hosting the reinforcement learning seminar and co-organizing the first deep reinforcement learning course in Poland. Prior to his academic career, he was leading a risk modeling team at KNF for multiple years, contributed to Solvency II directive, and was a member of EIOPA working groups; he is also a chartered actuary.
lkucinski@impan.pl

Piotr Miłoś is a professor at the Polish Academy of Sciences and a visiting professor at the University of Oxford. He specializes in machine learning and leads a research group focusing on reinforcement learning. Prof Milos has a strong mathematical background in probability theory. He actively works towards developing machine learning research in Poland. This includes hosting a reinforcement learning seminar and co-organize reinforcement learning course (first in Poland), and being a scientific advisor of the ML in PL association.
pmilos@impan.pl

Joanna Rączaszek-Leonardi is a full professor at the Department of Psychology, University of Warsaw where she heads the Human Interactivity and Language Lab. She has a PhD from the Center for Complex Systems and Brain Sciences, FAU, USA. Her background is in ecological psychology and dynamical systems. Her basic interest is in how symbolic forms of language emerge from and constrain the dynamics of human interactivity. Her theoretical works pertain to philosophy of language, cognitive sciences and embodied cognition, while she seeks empirical evidence in the areas of language development, experimental semiotics and computational modeling of language development and emergence.
raczasze@psych.uw.edu.pl

Publication history

Date received: 6 July 2021
Date accepted: 16 November 2021