

University of Warsaw
Faculty of Psychology

Tomasz Korbak

Student no. 331949

Emergent compositional communication via template transfer

Master's thesis in Cognitive science

Supervisors:

Prof. Joanna Rączaszek-Leonardi

Human Interactivity and Language Lab,
Faculty of Psychology, University of Warsaw

dr hab. Piotr Miłoś, prof. IMPAN

Institute of Mathematics,
Polish Academy of Sciences

December 2019

Supervisors' statement

Hereby I confirm that the presented thesis was prepared under my supervision and that it fulfils the requirements for the degree of Master of Cognitive Science.

Date

Supervisor's signature

Author's statement

Hereby I declare that the presented thesis was prepared by me and none of its contents was obtained by means that are against the law.

The thesis has never before been a subject of any procedure of obtaining an academic degree.

Moreover, I declare that the present version of the thesis is identical to the attached electronic version.

Date

Author's signature

Abstract

In this thesis, I explore a novel approach to achieving emergent compositional communication in multi-agent systems. I propose a training regime implementing template transfer, the idea of carrying over learned biases across contexts. In the presented method, a sender–receiver pair is first trained with a disentangled loss functions, and then the receiver is transferred to train a new sender with a standard loss. Unlike other methods (e.g. the obverter algorithm), template transfer approach does not require imposing inductive biases on the architecture of the agents. I 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, compositionality, Lewis signaling games, deep learning, disentangled representation, generalization

Thesis domain (Socrates-Erasmus subject area codes)

11.4 Artificial Intelligence

Tytuł pracy w języku polskim

Emergentna komunikacja kompozycyjna poprzez transfer szablonu

Contents

Acknowledgements	9
1. Introduction	11
2. Theoretical background	13
2.1. Lewis signaling games	13
2.1.1. Definitions	14
2.1.2. S -vector semantics	14
2.1.3. Template transfer and generalized signaling games	15
2.2. Hierarchies of communication protocols	16
2.3. Compositionality	18
3. Related work	21
4. Method	25
4.1. Experimental setup	25
4.1.1. Dataset	25
4.1.2. Object naming game	25
4.1.3. Derivation of the loss function for the object naming game	26
4.1.4. Reparametrization of the loss function for the object naming game	27
4.1.5. Architecture of the agents	28
4.2. Template transfer approach	29
5. Experiments and results	31
5.1. Measuring compositionality	31
5.2. Baselines	32
5.3. Results	32
5.4. Visualizing receiver’s activations	33
6. Discussion	37
7. Conclusions	41
Bibliography	41

List of Figures

4.1. Examples of images from the training dataset.	26
4.2. Object naming game	27
4.3. Template transfer consists of pre-training the receiver r on two games with disentangled losses L_1 and L_2 and transferring r to a new object naming game.	29
5.1. 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.	35
5.2. Receiver RNN's hidden states corresponding for each object type plotted on a 2d plane.	36

List of Tables

- 5.1. 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. 33
- 5.2. Two example communication protocols, one that emerged via the baseline architecture (5.2a), and one via template transfer (5.2a). Gray cells indicate objects not seen during training. In (5.2b), 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). 34

Acknowledgements

Significant parts of the thesis (especially chapters 1, 3-6) are based on an extended abstract accepted to the workshop Emergent Communication: Towards Natural Language at Conference on Neural Information Processing Systems (NeurIPS 2019) in Vancouver, Canada under the title Developmentally motivated emergence of compositional communication via template transfer and co-authored with Julian Zubek, SławoŃ Kuciński, Piotr Miłoś and Joanna Ręczaszek-Leonardi. My contribution to this research was supported by Ministry of Science and Higher Education (Poland) grant DI2015010945. Additionally, section 2.1 is based on my blog post titled Introduction to Lewis signaling games with Python and available under <https://tomekkorbak.com/2019/10/08/lewis-signaling-games/>.

The computational experiments presented in this thesis were run on the PLGRID infrastructure within the Prometheus cluster.

Chapter 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). One important feature of human languages and some animal communication systems is compositionality – there are complex signals constructed through the combination of signals. Compositionality is considered a key feature of general intelligence because it 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 compositionality requires strong inductive biases to be imposed on the agents (Kottur et al., 2017).

Contribution The contribution of this thesis is its demonstration 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 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) 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). We compare the template transfer approach with other method of achieving compositionality the oververter 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 oververter approach.

The structure of the thesis In chapter 1, I provide some background knowledge necessary to pose the problem of accounting for the emergence of compositional communication. This includes the formal model of (generalized) Lewis signaling games (section 2.1), the hierarchy of forms of communication as developed in Peircean semiotics (section 2.2) and the notion of compositionality itself (section 2.3). In chapter 3, I present a review of literature directly relevant for the studied approach. In chapter 4, I describe the task (including deriving its loss functions) and the setup used in experiments. Then I describe the particular implementation of template transfer considered in this thesis. In chapter 5, I describe metrics used in

the experiments, present the results of template transfer approach to the emergence of compositionality and compare them with some baseline approaches. In chapter 6, I develop evolutionary and developmental motivation behind my approach and argue it to be less cognitively demanding than other approaches. In chapter 7, I conclude by discussing some of the limitations of the present study and possible future work.

Chapter 2

Theoretical background

The goal of this chapter is to lay down the background necessary to pose the problem of accounting for the emergence of compositional communication. In the thesis, I 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. 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).

The structure of the chapter is as follows. First, in section 2.1 I explicate what I mean by communication in terms of a simple formal model – Lewis signaling games – and consider the problem of the emergence of a simple communication protocol from scratch in guessing game as an instance of Lewis signaling games. Next, I introduce a concept fundamental for the thesis – template transfer – and sketch how it accounts for the emergence of complex semiotic phenomena. Then, in section 2.2 I introduce another theoretical framework for tackling the problem of the emergence of complex communication out of simple communication – Peircean semiotics – and consider its modern extensions in cognitive anthropology and developmental psychology of language. Finally, in section 2.3 I introduce the notion of compositionality and its significance in cognitive science and artificial intelligence. I revisit some of the concepts from this chapter when discussing the implications of the empirical results of the thesis in chapter 6.

2.1. Lewis signaling games

Signaling transmits information, but it does far more than this. To see this we need to move further than the simple signaling games with one sender and one receiver. Signals operate in networks of senders and receivers at all levels of life. Information is transmitted, but it is also processed in various ways. Among other things, that is how we think just signals running around a very complicated signaling network. Very simple signaling systems should be able to learn to implement very simple information processing tasks by very simple means, and indeed they can.

Brian Skyrms

Lewis signaling games are a simple game-theoretic model of communication. The study of signaling games was pioneered by Lewis, who analyzed the emergence of social conventions in game-theoretic terms. While Lewis' motivation was philosophical (he was interested in

explicating and defending the neopositivist notion of truth by convention), a number of phenomena can be modeled as signaling games, including the labor market, sexual selection and animal communication (Skyrms, 2010).

2.1.1. Definitions

A Lewis signaling game demands a sender and a receiver to invent a communication protocol so that the receiver can act based on information only available to the sender and maximize reward for both of them. More formally, a Lewis signaling game consists of

1. a world, consisting of X (a set of states $x \in X$) and Y (a set of available actions $y \in Y$),
2. a set of available messages M ,
3. a sender $s : X \rightarrow M$ (a mapping from a world state to a message),
4. a receiver $r : M \rightarrow Y$ (a mapping from a message to an action), and
5. a loss function $L : X \times Y \rightarrow \mathbb{R}$ assigning each $(x; y)$ pair a scalar reward/penalty $l \in \mathbb{R}$.

Let us focus on a special case known as a guessing game in the language evolution literature (Cangelosi, 2001) (though not always formalized as a Lewis signaling game). We assume that the optimal action depends on the state of the world available only for the sender. In such a case, the sender is incentivized to transmit the information about the state to help the receiver make an informed decision. Furthermore, let us assume X and Y are equinumerous sets and are objects are mapped to correct actions via some bijective function $f : X \rightarrow Y$. Now let L be 0 if $r(s(x^{(i)})) = f(x^{(i)}) = y^{(i)}$ and some positive constant l otherwise. The goal of the guessing game is to find mappings s and r that recover f by minimizing L . A guessing game is solved when $L(x^{(i)}; y^{(i)}) = 0$ for all i (equivalently, when $s \circ r = f$). Intuitively, this reflects the following situation: the sender sees an object $x^{(i)}$ (one out of several available objects) and must communicate to the receiver which objects it sees so the receiver can to act to successfully indicate it among a set of distractors. If the receiver successfully indicates all objects in X based on sender's messages, the game is solved.

Let a communication protocol $(s; r)$ be a specific solution a particular guessing game². For each guessing game there are numerous possible communication protocols. The choice of a particular communication protocol is a social convention that the agents implicitly agree upon (as deviating from the protocol would increase the penalty for both of the agents).

2.1.2. S-vector semantics

A non-trivial semantics can be defined for a communication protocol, lending support to claims that the messages do indeed mean something and the emergence of meaning can be accounted for in terms of playing Lewis signaling games. More formally, a semantics for a given communication protocol is a function $[\cdot] : M \rightarrow M$ that assigns a unique, evaluable formal object $[m] \in M$ (a meaning) to each message $m \in M$ such that $[m]$ uniquely characterizes the content of the message m , e.g. determines in what circumstances m is to be sent (descriptive content) or what should be done upon receiving it (imperative content). Let us sketch an informal argument that for each communication protocol $(s; r)$ there exists a semantic $[\cdot]$.

Let us first assume that the sender and the receiver are stochastic, i.e. s and r are probability mass functions (parametrized by μ and ν , respectively) defining categorical distributions

¹There are more general formulations with information about the world available for both the sender and the receiver. I consider the special case where only the sender has information about the world because this is the setup used in my experiments.

²A communication protocol is equivalent to a signaling system in Lewis' terminology.

over messages and actions conditioned on states and messages, respectively:

$$m \sim s(m|x); \tag{2.1}$$

$$y \sim r(y|m); \tag{2.2}$$

Both distributions can be seen as policies parametrized by θ and ϕ , respectively. Let us assume θ and ϕ are initialized to parametrize uniform distributions. Solving a guessing game corresponds to the optimization problem

$$\arg \min_{\theta \in \Theta; \phi \in \Phi} \sum_{x \in X; y \in Y} E_{m \sim s(m|x^{(i)})} E_{y \sim r(y|m)} L(y^{(i)}; \phi); \tag{2.3}$$

Then, the emergence of meaning corresponds to the symmetry breaking of θ in the course of optimization: each message starts being sent in response to a progressively more restricted subset of object. The meaning of a message m in a guessing game can be defined in terms of how it affects $r(y|m)$, i.e. a change in $r(y)$ after receiving m gives rise to imperative content. The information conveyed by a message to the receiver is simply the point-wise mutual information between the message and the action, i.e. $\log \frac{r(y|m)}{r(y)}$. The information content of a message is a vector of point-wise mutual informations between the message and each of the actions available for the receiver. That is, for a game with m available actions for the receiver $(y^{(1)}; y^{(2)}; \dots; y^{(n)})$ the imperative meaning of a message m is

$$[m] = \langle \log \frac{r(y^{(1)}|m)}{r(y^{(1)})}; \log \frac{r(y^{(2)}|m)}{r(y^{(2)})}; \dots; \log \frac{r(y^{(n)}|m)}{r(y^{(n)})} \rangle; \tag{2.4}$$

Such a vector of log probability ratios is known as an s-vector. Intuitively, imperative content of messages describes how acquiring the message affects the behavior of the receiver. Note, however, that information content (a vector of log probability ratios) is a richer object than the quantity of information conveyed in a message (a scalar). How the message changes the distributions is more than just how much it changes the distribution. Imperative information content has a normative aspect as it describes the behavior expected of the receiver upon receiving the message.

The idea that the s-vector is a good candidate for a meaning was first proposed by Skyrms (2010) and has recently been defended by Isaac (2019) under the name of vector semantics. This approach is theoretically interesting because it has very few theoretical commitments: it only assumes the axioms of probability and the existence of a sender and a receiver. It also explicitly ties the meaning of a message to its use and how the meaning evolves, as well as to (Bayesian) inference involved in meaning comprehension and informed action. From a computational point of view, taking the set of available meanings to be a vector space is very in line with distributional approaches to semantics that underlie contemporary natural language processing (Mikolov et al., 2013).

2.1.3. Template transfer and 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 following classes of generalized signaling games:

1. a cue reading game is where $a_1 = s$ and $a_2 = r$, i.e. s is approximately fixed and the receiver takes advantage of the policy of the sender,
2. a sensory manipulation game is where $a_1 = r$ and $a_2 = s$, i.e. r 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).

Moreover, one might think of the ritualization of decisions as the glue that binds agents to form simple games from their basic decisions, then increasingly complex games from simple game (Barrett and Skyrms, 2017). In case when the policy of a_1 evolved as a solution to a previous signaling game g_0 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 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).

2.2. Hierarchies of communication protocols

Symbols grow. They come into being by development out of other signs.

Charles Sanders Peirce

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.

Importantly, being an icon, an index or a symbol is not an intrinsic property of a sign but a feature relative to an agent interpreting the sign. This follows from sign being a triadic relation between a sign vehicle, an object and an agent (interpretant in Peirce's terminology). Certain sign vehicles may be purely iconic for one agent while symbolic for another.

Deacon (1998) developed a cognitive anthropological interpretation of Peirce's semiotics and argued the linear order over three kinds of signs 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)³. Empirically, sensitivity to iconic reference can be found arbitrary 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).

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). Regarding development, the semiotic hierarchy of signs can be turned into typology of constraints language imposes on interacting agents. These constraints emerge in a structured social environment and depend on each other (in a sense that pre-existing constraints unleash novel forms of control) as part of a multithread complex process that at the same time maintains grounding of the system in which they are embedded and in which iconic and indexical grounding is progressively augmented or replaced by symbol-symbol relations (Riczaszek-Leonardi et al., 2018, p. 67).

The take-home message from Peirce and Deacon crucial for the approach presented in this thesis is that complex (symbolic or only compositional) communication protocols do not (usually) emerge from scratch. Instead, pre-existing, simpler communication protocols can be (and frequently are) employed as a scaffolding for the evolution of complex communication protocols.

³This is because symbolic reference between a symbol S and an object O is determined by a relationship S has with other symbols S^0, S^{00}, \dots , 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.

2.3. Compositionality

It is astonishing what language can do. With a few syllables it can express an incalculable number of thoughts, so that even a thought grasped by a human being for the very first time can be put into a form of words which will be understood by someone to whom the thought is entirely new.

Gottlob Frege

Compositionality is a property of certain communication protocols that the meaning of a complex expression is fully determined by its structure and the meanings of its constituents. More formally, consider a protocol consisting of a vocabulary $V = \{v_1; v_2; \dots; v_n\}$ and a binary syntactic composition operator f . Let us assume that a message has the form $m = v_i \ v_j$.⁴ Once again, let a semantics for a communication protocol be a function $\mathbb{f} : m \rightarrow \mathbb{F}([m])$ that assigns a unique, evaluable formal object $\mathbb{f}([m])$ to each message m such that $\mathbb{f}([m])$ uniquely characterizes the content of that message. A semantics is compositional if, for every v_i, v_j it holds that

$$\mathbb{f}([v_i \ v_j]) = f(\mathbb{f}([v_i]); \mathbb{f}([v_j])) \quad (2.5)$$

with f being a semantic composition function. Saying that a communication protocol is compositional is actually a short-cut for saying that agents use a compositional semantics for that protocol.

Compositionality is considered to be an essential feature of human languages Hockett (1960). Moreover, compositionality is assumed to be an important building block of general intelligence by being linked to productivity, systematicity and generalization. Productivity is the property that an unbounded number of meanings can be created using a finite number of primitive elements.⁵ This property is fundamental to several theories of universal grammar developed the generative approach in linguistics (Chomsky, 2015). Systematicity is the presence of definite and predictable patterns in the communication protocol which could potentially improve the learnability of the protocol. Systematicity can also be understood as a symmetry of a communication protocol with respect to composition (e.g. understanding the meaning of 'Eve loves Marry' entails understanding the meaning of 'Mary loves Eve'). Finally, generalization is the ability adapt to novel contexts. As such, it is central to machine learning and productivity and systematicity are frequently seen simply as means of improving generalization in some settings. This in particular involves compositional or combinatorial (zero-shot) generalization, i.e. adaptability to novel combinations of known elements (Lake, 2019; Hill et al., 2019).

Learning compositional representations is a long-standing challenge in artificial intelligence. Fodor and Pylyshyn (1988) have famously argued that a neural network must explicitly represent a compositional semantics in order understand a compositional communication protocol, because association (implemented as an affine transformation in a layer of a neural network) is not a structure-sensitive relation and structured representations could not be encoded by association.⁶ While Fodor and Pylyshyn's systematicity challenge was subsequently widely criticised and afforded book-length treatments (e.g. (Calvo and Symons, 2014)), a

⁴These simplifying assumptions reflect the experimental setup described in chapter 4 and linguistically correspond to (for example) a noun phrase composed of a noun and an adjective. Compositional semantics relevant for linguistics are recursively defined (i.e. m itself can be composed with some m^0) and allow for more sophisticated operators than f (Partee et al., 1993).

⁵This assumes a recursively defined semantic composition function f .

⁶Note that, for similar reasons, s-vector semantics described in section 2.1 also fails to be compositional because one cannot infer the joint distribution $p(v_i; v_j)$ from marginal distributions $p(v_i)$ and $p(v_j)$.

resurgence in neural network research in 2010s invigorated interest in the problem (Lake et al., 2016). Modern research assumes compositionality to be a graded notion and studies the impact of various inductive biases on compositional generalization as a proxy metric for compositionality of representations (Hill et al., 2019). Literature directly relevant to the problem of compositionality in emergent communication discussed in the next chapter.

Chapter 3

Related work

In this chapter, I review recent empirical work on deep neural network-based models of the emergence of compositionality in signaling games. First, I discuss various mechanisms known to incentivize compositional communication. Then, within this research tradition, I state the problem to be solved in the thesis.

Towards deep learning Computational models of 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 policies of agents and r . With the rise of deep learning (Goodfellow et al., 2016), deep neural networks started being used to implement policies with parameters optimized via gradient descent implemented using the backpropagation algorithm (Rumelhart et al., 1986). In a typical setting, learning boils down to minimizing $r ; L(y; r (s (x)))$ for L , x and y defined as in chapter 2. 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 constraints 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 placing 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 *obvertō* to turn towards) algorithm (Batali, 1998; Oliphant and Batali, 1997) 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 and a' (with policies parametrized by θ_a and $\theta_{a'}$) exchange the roles of the sender and the receiver. If an agent is the receiver ($\hat{a} = r$), it behaves as in the object naming game. If an agent is the sender ($\hat{a} = s$), it sends message that would have produced the optimal response (to the best of a' 's knowledge) if a' had received such a message as a receiver. More formally, a sends a message $m = \arg \max_{m \in \mathcal{M}} (y_{c|jm}^0)$, assuming a correct action y_c is known or can be predicted by a' . 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 (y_{jm}) based on own policy (y_{jm}).

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 must share an identical architecture 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 next symbol maximizing $r(y_{c|jm_{1:t}})$), producing a message requires $\mathcal{O}(jVjT)$ queries to the model of the receiver (where jVj is vocabulary size and T is maximum message length).

Population-based training A different family of approaches tries to incentivize 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 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.

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 in inducing 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 *love* and *Mary* might reuse

it together with new words for novel elements of the event. The experiments conducted by De Beule and Bergen 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.

Grounded and situated approaches Language grounding, situatedness and embodiment are emerging trends in recent research in both cognitive science and artificial intelligence (Pfeifer and Scheier, 1999; Smith and Gasser, 2005; Clark, 2016), posing questions such as the emergence of compositional communication in more psychologically realistic settings. Interestingly, such settings may yield novel patterns of behavior not appearing in simpler, toy environments. The departure from pre-processed, symbolic input (e.g. encoding an object as a one-hot vector) in favor of raw visual input (i.e. tensors of RGB pixels) requires the sender to deal with entangled input and decompose it into relevant factors of variation (e.g. shape and color) on its own, yielding less compositional communication protocols, on average (Lazaridou et al., 2018). Communication in an embodied and situated settings was studied by Mordatch and Abbeel (2017). Their environment is a 2d continuous space populated by agents and colored landmarks; agents learn to communicate to each other about actions that they must undertake (e.g. go to) and object that actions must target (e.g. blue), developing a communication protocol with compositional verb noun structure. The ability of the agent to move freely contrasts with dominant settings with predefined roles of a sender and a receiver, where the sender passively observes a stimulus (independently of any actions it can undertake) and the receiver is waiting for a message to undertake an action. This enabled non-verbal communication (the transfer of valuable information outside of a predefined communication channel), e.g. by gaze direction or physical movement. Moreover, according to Mordatch and Abbeel, physical environment considerations play a part in the syntactic structure. The action type verb GOTO is uttered first because actions take time to accomplish in the grounded environment. When the agent receives GOTO symbol it starts moving toward the centroid of all the landmarks (to be equidistant from all of them) and then moves towards the specific landmark when it receives its color identity (Mordatch and Abbeel, 2017).

Situatedness was also shown to have a positive influence on compositional generalization. Hill et al. (2019) demonstrate that (i) egocentric perspective (as opposed to third-person view), (ii) 3d simulated environments (as opposed to 2d grid worlds) and (iii) active perception (as opposed to perceiving the stimuli statically) incentivize the agent to factorise experience and behaviors into reusable chunks, leading to better compositional generalization as shown in experiments isolating factors (i)-(iii). These results provide explanation of previous results, showing low systematic generalization in recurrent neural networks trained on symbolically encoded stimuli from a single modality (Lake and Baroni, 2017). The authors hypothesize that this effect arises because richer experience is a form of implicit data augmentation, which suggests that the human capacity to exploit the compositionality of the world, when learning to generalize in systematic ways, might be replicated in artificial neural networks if those networks are afforded access to a rich, interactive, multimodal stream of stimuli that better matches the experience of an embodied human learner (Hill et al., 2019).

Motivation for the presented approach The problem addressed in this thesis is to propose and evaluate an alternative, novel approach to the emergence of compositional

communication. Following the literature, we will pose the problem as anaming game a variation of a guessing game, where the sender sees an object with two independent factors of variation (shape and color) and the receiver must, independently, indicate both of these factors. While this treatment of compositionality is simplistic and amounts to what Steinert-Threlkeld (forthcoming) calls generalized conjunction or trivial compositionality, even trivially compositional communication protocols remain difficult to learn (as seen in previous paragraphs) and provide a minimal model of the phenomenon. The objects will be presented to the sender as raw pixel data, which is motivated by (relative) biological plausibility of this setting. There will be a pre-defined communication channel: a set of 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-symbols. Both agents will be implemented as recurrent neural networks with time-steps corresponding to subsequent symbols. The implementations details of the experimental setup are presented in the next chapter.

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, I will investigate whether template transfer (as described in section 2.1) can be employed as a way of achieving compositional communication in an object naming game.

Chapter 4

Method

The goal of this chapter is to describe the experimental setup in more detail, derive the specific loss function used in experiments and present the template transfer approach – the main contribution of the thesis.

4.1. Experimental setup

4.1.1. Dataset

We conduct our experiments on a dataset consisting of 2500 images of colored three-dimensional objects. Each image has dimensions of $28 \times 28 \times 3$ 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 engine¹ 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 shape and color, i.e. the test set is composed of blue boxes, cyan spheres, gray cylinders, green tori and magenta ellipsoids. Example images from the training set are shown on figure 4.1.

4.1.2. Object naming game

Object naming games are Lewis signaling games, which extend the guessing game framework (as described in section 2.1) to a setting where the loss function can be decomposed into a sum of two loss functions L_1 and 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:

$$L(\theta; \mathcal{D}) := \mathbb{E}_{x, y_c, y_s \sim \mathcal{D}} \mathbb{E}_{m \sim s(\cdot|x)} [\log r(y_c; y_s|m)]; \quad (4.1)$$

where s is the policy of the sender (i.e. $s(m|x)$ is the probability of sending message m when observing image x), r is the policy of the receiver (i.e. $r(y_c; y_s|m)$ is the probability of taking actions $y_c; y_s$ after receiving message m). \mathcal{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 ϕ are learnable parameters of the policies. For more details, see Algorithm 1 and Figure 4.2.

¹The dataset was generated using code available from <https://github.com/benbogin/obverter>.

Figure 4.1: Examples of images from the training dataset.

Algorithm 1 Training loop for the object naming game

```

1: Initialize sender  $s$ , receiver  $r$ , and training set  $D$ 
2: for  $x; y_c; y_s \in D$  do
3:    $m = s(x)$ 
4:    $\hat{y}_c; \hat{y}_s = r(m)$ 
5:    $L = -\log_{\text{likelihood}}(y_c, \hat{y}_c) - \log_{\text{likelihood}}(y_s, \hat{y}_s)$ 
6:   optimize( $L$ ;  $\theta$ )

```

4.1.3. Derivation of the loss function for the object naming game

For completeness, let us derive (4.1) as a maximum likelihood estimator of a solution to object naming game given the topology of the graphical model embodied by the sender receiver pair mediated by a discrete latent variable (the message). Let us assume that we are given a dataset

$$D = \left\{ (x^{(i)}; y_c^{(i)}; y_s^{(i)}) \right\}_{i=1}^n; \quad (4.2)$$

where entries are independent and identically distributed, $x^{(i)}$ is an object (encoded as an RGB image) and $y_c^{(i)}, y_s^{(i)}$ are (ground truth) labels for $x^{(i)}$ (encoded as one-hot vectors). Let us further assume that each object comes from a distribution $(X; Y_s; Y_c; M)$, where Y_s and Y_c are two labels for image X and M is a latent variable. The goal of a naming game is minimizing the negative log likelihood of ground truth labels given the image

$$E_{(x; y_c; y_s) \in D} [-\log p(y_c; y_s | x)]; \quad (4.3)$$

Figure 4.2: Object naming game

Assuming that $(Y_c; Y_s)$ and X are conditionally independent given M , i.e. $p(y_s; y_c | x; m) = p(y_s; y_c | m)$, we have

$$\log p(y_c; y_s | x) = \log \prod_m p(y_c; y_s | m) p(m | x) = \mathbb{E}_{m \sim p(\cdot | x)} [\log p(y_c; y_s | m)] \quad (4.4)$$

where the last inequality follows from Jensen's inequality. We will be optimizing the lower bound; hence, our surrogate loss function is

$$\mathbb{E}_{(x; y_c; y_s) \sim D} \mathbb{E}_{m \sim p(\cdot | x)} [\log p(y_c; y_s | m)]: \quad (4.5)$$

By substitution of the first and second probability mass function with sender's policy s and receiver's policy r , respectively, we obtain (4.1).

4.1.4. Reparametrization of the loss function for the object naming game

In our experiment we will leverage gradient descent optimization to jointly train the sender and the receiver to minimize \mathcal{L} . One problem for (4.1) as a loss function is that it is not differentiable with respect to θ due to non-differentiable sampling from the expectation $\mathbb{E}_{m \sim p(\cdot | x)}$. Moreover, m over which the expectation is taken is discrete, which is a crucial assumption from a linguistic point of view. This prevents direct use of reparametrization (Kingma and Welling, 2013) moving θ from the distribution to inside the expectation or vice versa because a probability mass function of a categorical distribution is not reparametrizable (Schulman et al., 2015). Instead, we will estimate the true gradient $\mathbb{E}_{m \sim p(\cdot | x)} [\log r(y_c; y_s | m)]$ by first relaxing the categorical distribution $p(M | X)$ and, subsequently, by reparametrizing its probability density function.

Let us assume that the policy of the sender encodes a categorical distribution over symbols $\{v_1; v_2; \dots; v_n\}$ that form a message $m = (v^0; v^0)$ for $v^0, v^0 \in V$. At each time-step, the RNN (parametrized by θ) predicts the probabilities $p_1; p_2; \dots; p_n$ for generating a symbols $v_1; v_2; \dots; v_n$. Instead of sampling v_i (represented as a one-hot vector) directly from this distribution, we obtain a sample from the corresponding Gumbel-softmax (Jang et al., 2016; Maddison et al., 2016) distribution given by

$$v_i = \frac{\exp((\log p_i + g_i) / \tau)}{\sum_j \exp((\log p_j + g_j) / \tau)} \quad (4.6)$$

where τ is a temperature parameter controlling the degree of relaxation and g_k (for $k \in [1, n]$) are samples from a standard Gumbel distribution, i.e. $g_k = -\log(-\log(u))$, where u is a sample from the uniform distribution from 0 to 1.

Then, the relaxed representation of the sample has the form of

$$\mathbf{v} = (v_1; \dots; v_n) \quad (4.7)$$

Gumbel-softmax samples approximate one-hot samples from the categorical distribution given by parameters $p_1; p_2; \dots; p_n$. By plugging in $\mathbf{m} = (\hat{v}_1^0; \hat{v}_2^0)$, this allows us to rewrite (4.1) to the form

$$L(\mathbf{v}; \mathbf{s}) := \mathbb{E}_{\mathbf{x}; y_c; y_s \sim D} \mathbb{E}_{u \sim U(0,1)} [\log r(y_c; y_s | \mathbf{s}(\mathbf{x}; u))]: \quad (4.8)$$

Crucially, \mathbf{s} is now a deterministic function of the object \mathbf{x} and a random sample u from the uniform distribution $U(0, 1)$, which makes L differentiable with respect to \mathbf{v} .

4.1.5. Architecture of the agents

General setup Both the sender and the receiver are implemented as recurrent neural networks. The sender is equipped with a pre-trained convolutional neural network to process visual input. After observing the object, the sender generates a sequence of discrete messages sampled from a closed vocabulary of 10 symbols.

All experiments reported in this thesis are implemented using PyTorch (Paszke et al., 2017) and EGG (Kharitonov et al., 2019). The code is publicly available²

Vision module We pre-train a simple convolutional neural network on the training subset of our dataset to predict colors and shapes. The network is composed of two layers of filters (20 and 50 filters with kernel size 5x5 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 cost 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 that is a trainable parameter as well). A symbol is sampled from this distribution at each time-step. 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.

Training 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} .

²The code was released on GitHub under <https://github.com/tomekkorbak/compositional-communication-via-template-transfer>.

4.2. 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:

$$L_1(\theta; \mu) := E_{(x; y_c) \sim D} E_{m \sim s_1(j|x)} [\log r(y_c | j, m)]; \quad (4.9)$$

$$L_2(\theta; \mu) := E_{(x; y_s) \sim D} E_{m \sim s_2(j|x)} [\log r(y_s | j, m)]; \quad (4.10)$$

where $r(y_c | j, m)$ is the marginalization of $r(y_c; y_s | j, m)$, viz. $r(y_c | j, m) := \int_{y_s} r(x_c; x_s | j, m)$. Analogously, one can define $r(y_s | j, m) := \int_{y_c} r(x_c; x_s | j, m)$. Crucially, as far as Y_c is conditionally independent from Y_s given X , we have

$$L(\theta; \mu) = L_1(\theta; \mu) + L_2(\theta; \mu); \quad (4.11)$$

The loss functions L_1 and L_2 are optimized simultaneously (crucially with the shared parameters 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 receiver. 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. 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 4.3 and Algorithm 2 for more details.

Figure 4.3: Template transfer consists of pre-training the receiver on two games with disentangled losses L_1 and L_2 and transferring r to a new object naming game.

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 chapter 5 indicate that two-phase learning is a sufficient incentive for compositionality to emerge.

The presented approach instantiates both template transfer and modular composition as described in section 2.1. To simplify the notation, let us assume for a moment that s_1 and

Algorithm 2 Template transfer approach

```

1: Initialize senders  $s_1, s_2, s$ ; receiver  $r$ , and training set  $D$ 
2: for  $x; x_c; x_s \in D$  do
3:    $m_1 = s_1(x)$  . Color naming game
4:    $m^0$  vocabulary
5:    $\hat{x}_c; \hat{x}_s = r([m_1; m^0])$ 
6:    $L_1 = -\log\_likelihood(x_c, \hat{x}_c)$ 
7:    $m_2 = s_2(x)$  . Shape naming game
8:    $m^{00}$  vocabulary
9:    $\hat{x}_c; \hat{x}_s = r([m^{00}; m_2])$ 
10:   $L_2 = -\log\_likelihood(x_s, \hat{x}_s)$ 
11:  optimize( $(L_1(\cdot; \cdot) + L_2(\cdot; \cdot))$ )
12: for  $x; x_c; x_s \in X$  do
13:   $m = s(x)$  . Object naming game
14:   $\hat{x}_c; \hat{x}_s = r(m)$ 
15:   $L = -\log\_likelihood(x_c, \hat{x}_c) - \log\_likelihood(x_s, \hat{x}_s)$ 
16:  optimize( $L(\cdot; \cdot)$ )

```

r are deterministic functions $s(x) := \arg \max_m s(m|x)$ and $r(m) := \arg \max_y r(y|m)$ (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: $g_a(r(s_1(x)); r(s_2(x)))$ such that the loss for g_a is $L_1 + L_2$. The presented approach also instantiates template transfer from the composite game g_a to the object naming game. In the latter game, the new sender s takes advantage of the biases in the receiver due to playing the composite game. The significance of these interpretations is further discussed in chapter 6.

Chapter 5

Experiments and results

In this chapter, I attempt to measure how much does the template transfer approach influence the degree of compositionality of a communication protocol as compared to three baseline approaches (random agents, the same architecture without pre-training and the oververter approach). The compositionality is measured using three metrics: test accuracy, context independence and topographical similarity, which will be described in the first section. Finally, I also try to provide an attempt at explaining how template transfer affects 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.

5.1. 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 $s(x) := \arg \max_m s(m|x)$, 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.

Contextual 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 K by 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. Further, let $v^k := \arg \max_v p(k|v)$. The context independence metric is defined as $E(p(v^k|k) - p(k|v^k))$; 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 $t \in \{0; 1; 2\}$. 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 Spearman 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 stimuli and a neural activity evoked by the stimuli (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.

5.2. 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.

5.3. 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); the results are presented in Table 5.1. Template transfer clearly leads to highly compositional communication protocols. While all methods struggled to generalize to unseen objects, template transfer was the most successful. For examples

Algorithm 3 Obverter

```
1: Initialize agents  $a_1; a_2$ , visual module  $v$ , training set  $D$ 
2: Initialize the set  $M$  of all possible messages
3: for  $x; x_c; x_s \in D$  do
4:    $s; r \in \{a_1; a_2\}$  . Randomly assigning the roles of sender and receiver
5:    $m = \arg \min_{m \in M} \text{evaluate\_message}(s, m)$ 
6:    $\hat{y}_c; \hat{y}_s = r(m)$ 
7:    $L = -\log_{\text{likelihood}}(x_c, \hat{y}_c) - \log_{\text{likelihood}}(x_s, \hat{y}_s)$ 
8:   optimize( $L$ )
9: procedure evaluate_message (model,  $m$ )
10:   $x_c; x_s = v(x)$  . Using visual classifier predictions as a proxy for ground truth labels
11:   $\hat{y}_c; \hat{y}_s = \text{model}(m)$ 
12:   $L^0 = -\log_{\text{likelihood}}(x_c, \hat{y}_c) - \log_{\text{likelihood}}(x_s, \hat{y}_s)$ 
13:  return  $L^0$ 
```

of communication protocols representative of the experiments conducted, see Table 5.2 and Figure 5.1.

Table 5.1: 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.

Model	Accuracy			CI	Topo
	Train (both)	Test (both)	Test (avg)		
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)
TT (ours)	1 (0)	0.48 (0.10)	0.74 (0.06)	0.18 (0.01)	0.85 (0.03)

5.4. Visualizing receiver's activations

Recall that the receiver r 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 $\text{proj}_{p_1, p_2} h_m$ of each hidden state onto two principal components p_1 and p_2 .¹ The scatter plots visualizing the RNN hidden states for the baseline architecture and template transfer are shown on Figure 5.2.

¹It is common in the literature to use PCA as a method for visualizing hidden states of RNNs, see also Yamashita and Tani (2008)

Table 5.2: Two example communication protocols, one that emerged via the baseline architecture (5.2a), and one via template transfer (5.2a). Gray cells indicate objects not seen during training. In (5.2b), 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	1 0	4 5	1 0	4 5	5 0
cyan	9 0	4 0	3 0	4 0	7 0
gray	3 5	6 5	3 2	6 5	5 3
green	0 0	7 6	3 0	6 0	7 6
magenta	1 5	5 5	1 2	1 5	5 2

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

	box	sphere	cylinder	torus	ellipsoid
blue	1 8	1 9	1 5	1 6	1 4
cyan	4 8	4 9	4 5	4 6	4 4
gray	6 8	6 9	6 5	6 6	6 9
green	9 8	9 9	9 5	9 6	9 4
magenta	8 8	8 9	8 5	8 8	8 4

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).

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

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

Figure 5.1: 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.

(a) No template transfer

(b) Template transfer

(c) Legend

Figure 5.2: Receiver RNN's hidden states corresponding for each object type plotted on a 2d plane.

Chapter 6

Discussion

In this chapter, I discuss the implications of the presented approach, focusing on the following points: (i) how compositionality can be understood game-theoretically, (ii) how the results corroborate Deacon account of reference, (iii) that compositionality may be less cognitively demanding than previously thought, and (iv) that the presented approach is developmentally plausible to a certain extent.

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) as described in section 2.1. The subsequent discussion focused 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 (compare with f in section 2.3), and modular composition, an operation over games yielding composite games. The assumption central to the presented approach is that a game (such as the object naming game) can be reformulated as a modular composition of two simpler games: $g_a(r(s_1(x)); r(s_2(x)))$ 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_1 and s_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. From an evolutionary point of view, one can 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.

Semiotic interpretation The presented 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 finally (3) learning a compositional communication protocol. This maps into the hierarchy of different forms of reference similar to the one described in section 2.2: (1) iconic reference, (2) indexical reference and (3) complex indexical reference¹ (Peirce, 1998; Deacon, 1998). The results corroborate the

¹I 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.

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 of 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_1 and r and between s_2 and r . 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_1 and s_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 to communicate compositionally.

Cognitive interpretation Some of the existing methods of inducing compositionality (e.g. the oververter approach) focus on imposing strong inductive biases on the architecture of the agents (cf. chapter 3). For instance, the oververter 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: that the loss function can be decomposed into two disentangled loss functions (as in the case of decomposing L into L_1 and L_2 in (4.9)-(4.10). (Note that there is no need for the input to be disentangled.) The fact that template transfer can outperform the oververter approach on all compositionality metrics lends support to the claim that the cognitive requirements for developing a compositional communication protocol are quite modest.

One may argue that template transfer obscures 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 (Riczaszek-Leonardi, 2012). Assuming this picture, language is an activity harnessed by constraints that are physical structures, 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 took advantage of the solution to the coordination problem developed jointly by s_1 , s_2 , and inherited it implicitly, never interacting with s_1 and s_2 . 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; Ryczaszek-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.

Chapter 7

Conclusions

The goal of the thesis 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.

The presented thesis 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_1 and s_2 . 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, where more complex (e.g. recursively nested) compositionality 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.

Bibliography

- Barrett, J., Skyrms, B., and Cochran, C. (2018). Hierarchical Models for the Evolution of Compositional Language.
- Barrett, J. A. and Skyrms, B. (2017). Self-assembling Games. The British Journal for the Philosophy of Science, 68(2):329-353.
- Barrett, J. A., Skyrms, B., and Mohseni, A. (2019). Self-Assembling Networks. The British Journal for the Philosophy of Science, 70(1):301-325.
- 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.
- Brighton, H. (2002). Compositional Syntax From Cultural Transmission. Artificial Life, 8(1):25-54.
- Brighton, H. and Kirby, S. (2006). Understanding Linguistic Evolution by Visualizing the Emergence of Topographic Mappings. Artificial Life, 12(2):229-242.
- 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 edition.
- Calvo, P. and Symons, J., editors (2014). The architecture of cognition: rethinking Fodor and Pylyshyn's systematicity challenge. The MIT Press, Cambridge, Massachusetts.
- Cangelosi, A. (2001). Evolution of communication and language using signals, symbols, and words. IEEE Transactions on Evolutionary Computation, 5(2):93-101.
- Choi, E., Lazaridou, A., and de Freitas, N. (2018). Compositional Obverter Communication Learning From Raw Visual Input. ICLR 2018. arXiv: 1804.02341.
- Chomsky, N. (2015). Syntactic structures. Martino, Mansfield Centre, Conn. OCLC: 934673149.
- Clark, A. (2016). Surging uncertainty: prediction, action, and the embodied mind. Oxford University Press, Oxford ; New York.

- 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.
- 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.
- De Beule, J. and Bergen, B. K. (2006). On the emergence of compositionality. The Evolution of Language, pages 35-42, Rome, Italy. World scientific.
- Deacon, T. W. (1998). The symbolic species: the co-evolution of language and the brain. Norton, New York, NY, norton paperback edition. OCLC: 254499872.
- Fodor, J. A. and Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: a critical analysis. Cognition, 28(1-2):3-71.
- 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.
- 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.
- Hill, F., Lampinen, A., Schneider, R., Clark, S., Botvinick, M., McClelland, J. L., and Santoro, A. (2019). Emergent Systematic Generalization in a Situated Agent.
- Hockett, C. F. (1960). The origin of speech. Scientific American, 203:89-96.
- Isaac, A. M. C. (2019). The Semantics Latent in Shannon Information. The British Journal for the Philosophy of Science, 70(1):103-125.
- Jang, E., Gu, S., and Poole, B. (2016). Categorical Reparameterization with Gumbel-Softmax. arXiv:1611.01144[cs, stat]. arXiv: 1611.01144.
- Jaquez, 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., 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.
- Kingma, D. P. and Ba, J. (2014). Adam: A Method for Stochastic Optimization. arXiv:1412.6980[cs]. arXiv: 1412.6980.
- Kingma, D. P. and Welling, M. (2013). Auto-Encoding Variational Bayes. CoRR, abs/1312.6114.

- 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.
- 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.
- 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.
- Kriegeskorte, N. (2008). Representational similarity analysis connecting the branches of systems neuroscience. Frontiers in Systems Neuroscience.
- Lake, B. M. (2019). Compositional generalization through meta sequence-to-sequence learning. CoRR, abs/1906.05381.
- Lake, B. M. and Baroni, M. (2017). Generalization without Systematicity: On the Compositional Skills of Sequence-to-Sequence Recurrent Networks. ICML.
- 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.
- Levenshtein, V. I. (1966). Binary Codes Capable of Correcting Deletions, Insertions and Reversals. Soviet Physics Doklady, 10:707.
- 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.
- Maddison, C. J., Mnih, A., and Teh, Y. W. (2016). The Concrete Distribution: A Continuous Relaxation of Discrete Random Variables. arXiv:1611.00712[cs, stat]. arXiv: 1611.00712.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. (2013). Distributed Representations of Words and Phrases and their Compositionality. In Burges, C. J. C., Bottou, L., Welling, M., Ghahramani, Z., and Weinberger, K. Q., editors, Advances in Neural Information Processing Systems 26, pages 3111-3119. Curran Associates, Inc.
- Mordatch, I. and Abbeel, P. (2017). Emergence of Grounded Compositional Language in Multi-Agent Populations. In AAAI.
- 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.

