
baller2vec++: A Look-Ahead Multi-Entity Transformer For Modeling Coordinated Agents

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 In many multi-agent spatiotemporal systems, agents operate under the influence of
2 shared, *unobserved* variables (e.g., the play a team is executing in a game of bas-
3 ketball). As a result, the trajectories of the agents are often statistically *dependent*
4 at any given time step; however, almost universally, multi-agent models implicitly
5 assume the agents’ trajectories are statistically *independent* at each time step. In
6 this paper, we introduce `baller2vec++`¹, a multi-entity Transformer that can effec-
7 tively model coordinated agents. Specifically, `baller2vec++` applies a specially
8 designed self-attention mask to a *mixture* of location *and* “look-ahead” trajectory
9 sequences to learn the distributions of statistically dependent agent trajectories. We
10 show that, unlike `baller2vec` (`baller2vec++`’s predecessor), `baller2vec++`
11 can learn to emulate the behavior of perfectly coordinated agents in a simulated
12 toy dataset. Additionally, when modeling the trajectories of professional basketball
13 players, `baller2vec++` outperforms `baller2vec` by a wide margin.

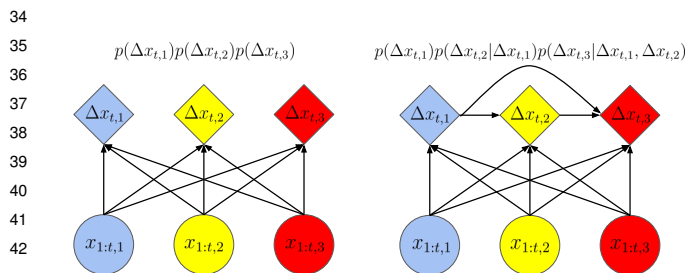
14 1 Introduction and Related Work

15 Whether it is a team executing a play in a game of basketball, a family navigating to an attraction in a
16 theme park, or friends posting about a birthday party on a social media platform, humans frequently
17 coordinate their behavior in response to shared information. When this coordinating information is
18 unobserved (which is often the case in many machine learning datasets), the individuals’ observed
19 behaviors become correlated, i.e., the behavior of one individual at a specific moment contains
20 information about the behavior of another individual *at the same time*. In the context of modeling
21 agent trajectories in multi-agent spatiotemporal systems, this property translates to the trajectories
22 being statistically dependent at each time step. However, nearly all multi-agent spatiotemporal models
23 (e.g., [1–6]) implicitly (through their loss functions) assume the trajectories of the agents at each time
24 step are statistically *independent* given the agents’ previous locations (Figure 1).

25 Zhan et al. [7] explicitly focused on modeling coordinated multi-agent trajectories, using “macro-
26 intents” [8] that are shared across agents to do so. The macro-intents are generated from a
27 separately trained recurrent neural network (RNN) that learns to predict a future, coarse, “sta-
28 tionary” location for each agent at each time step. The macro-intents for all of the agents at
29 a specific time step are concatenated together to form a single, shared, macro-intent variable,
30 which is then provided as input to the trajectories-generating model at that time step. However,
31 similar to the previously mentioned multi-agent trajectory models, the macro-intents model im-

¹All data and code for the paper are available at: <anonymized>.

32 plicitly assumes the macro-intents for the agents at each time step are statistically independent,
 33 i.e., the macro-intent for one agent does not depend on the macro-intents of the other agents.²



43
 44 **Figure 1: Left:** most multi-agent systems implicitly assume the trajectories of the agents at each time step
 45 $(\Delta x_{t,k})$ are conditionally independent given the agents’ previous locations $(x_{1:t,k})$. **Right:** however, the various de-
 46 compositions of the joint probability of the trajectories, e.g., $p(\Delta x_{t,1})p(\Delta x_{t,2}|\Delta x_{t,1})p(\Delta x_{t,3}|\Delta x_{t,1}, \Delta x_{t,2})$ (note,
 47 we omit the conditional $x_{1:t,k}$ terms for brevity), suggest more complex statistical dependencies between the agents’
 48 trajectories can exist (i.e., the independence assumption is an extremely strong one). Indeed, there are often shared un-
 49 observed variables influencing the spatiotemporal behaviors of agents—such as the play that the players on a basketball
 50 team are executing, or events occurring in a pedestrian environment—which suggests statistical dependencies between
 51 the agents’ trajectories are likely.

52
 53
 54
 55
 56
 57
 58
 59
 60
 61 estingly, adding the global discriminator to a baseline model only improved the model’s performance
 62 for one out of six pedestrian datasets.

63 In this paper, we describe a novel multi-agent spatiotemporal model that integrates information about
 64 *concurrent* actions of agents to predict statistically dependent distributions of trajectories. Specifically,
 65 we extend the recently introduced multi-entity Transformer `baller2vec` [6] by: (1) augmenting its
 66 input with a parallel sequence of “look-ahead” agent trajectories and (2) using a specially designed
 67 self-attention mask, which allows our model to exploit the chain rule of probability (Section 3). We
 68 find that:

- 69 1. `baller2vec++` is an effective learning algorithm for modeling coordinated agents. Unlike
 70 `baller2vec`, `baller2vec++` can learn to emulate perfectly coordinated agents from a
 71 simulated toy dataset (Section 5.1). Further, `baller2vec++` outperforms `baller2vec` by
 72 a wide margin (8.9%) when modeling the trajectories of professional basketball players
 73 (Section 5.1).
- 74 2. `baller2vec++` makes better predictions when conditioned on concurrent trajectory infor-
 75 mation from other agents, supporting our proposition that the commonly used independence
 76 assumption for agent trajectories is overly strong (Section 5.2).
- 77 3. Lastly, the joint probability assigned to a sequence by `baller2vec++` is approximately
 78 permutation invariant with respect to the order of the agents, i.e., `baller2vec++` respects
 79 the properties of the chain rule (Section 5.3).

²See the authors’ implementation here: https://github.com/ezhan94/multiagent-programmatic-supervision/blob/a1d9152d4c8a287474953cba093c28fef2a05979/models/macro_vrnn.py#L101.

80 2 Background

81 2.1 Multi-agent trajectory modelling

82 Our problem description closely follows Alcorn and Nguyen [6], whom we quote here:

83 Let $A = \{1, 2, \dots, B\}$ be a set indexing B agents and $P = \{p_1, p_2, \dots, p_K\} \subset$
 84 A be the K agents involved in a particular sequence. [Let] $C_t =$
 85 $\{(x_{t,1}, y_{t,1}), (x_{t,2}, y_{t,2}), \dots, (x_{t,K}, y_{t,K})\}$ [be] an unordered set of K coordinate
 86 pairs such that $(x_{t,k}, y_{t,k})$ are the coordinates for agent p_k at time step t . The
 87 ordered sequence of sets of coordinates $\mathcal{C} = (C_1, C_2, \dots, C_T)$, together with P ,
 88 thus defines the trajectories for the K agents over T time steps.

89 In multi-agent trajectory modeling, the goal is to model a joint probability of the form:

$$p(\Delta x_{t,1}, \Delta x_{t,2}, \dots, \Delta x_{t,K} | x_{1:t,1}, x_{1:t,2}, \dots, x_{1:t,K})$$

90 i.e., the joint probability of the K agents’ trajectories $\Delta x_{t,k}$ at time step t given the agents’ location
 91 histories $x_{1:t,k}$. We note here that the common practice of *simultaneously* predicting the trajectories
 92 for all of the agents at a specific time step is not required by theory. Using the chain rule of probability,
 93 the joint probability of the agents’ trajectories can be factorized as, e.g.:

$$p(\Delta x_{t,1}, \Delta x_{t,2}, \dots, \Delta x_{t,K}) = p(\Delta x_{t,1})p(\Delta x_{t,2} | \Delta x_{t,1}) \dots p(\Delta x_{t,K} | \Delta x_{t,1}, \Delta x_{t,2}, \dots, \Delta x_{t,K-1})$$

94 where we omit the conditional historical trajectories for brevity. As a result, it is perfectly acceptable to
 95 generate trajectories agent-wise, using the previously generated trajectories as additional conditioning
 96 information when generating the trajectories for later agents (see Figure 2).




Figure 2: At inference time, a model is *not required* to *simultaneously* generate the trajectories for all of the agents at a specific time step. An alternative strategy is to allow the model to generate the agents’ trajectories one at a time, and let the model use the previously generated trajectories to inform the trajectories it generates for the remaining agents.

97 2.2 baller2vec is a (conditional) generative model.

98 baller2vec is a recently described *multi-entity* Transformer that can model sequences of *sets*
 99 (the underlying data structure for multi-agent spatiotemporal systems), as opposed to sequences
 100 of individual inputs (like words in a sentence). When used to model the game of basketball, the
 101 input at each time step for baller2vec is a set of feature vectors where each feature vector contains
 102 information about the identity and location of a player on the court. baller2vec maps each input
 103 feature vector to an output feature vector, which is then used to “classify” the binned trajectory for
 104 that specific player at that specific time step.

105 Here, we provide a probabilistic interpretation of baller2vec, which establishes the theoret-
 106 ical grounds for using the chain rule to generate trajectories agent-wise at each time step in
 107 baller2vec++. Without loss of generality, we only consider one-dimensional trajectories for
 108 a single agent here. To briefly summarize, the outputs of the softmax function over the n binned

109 trajectories in `baller2vec` can be interpreted as mixture proportions for a mixture of uniform distri-
 110 butions with predetermined bounds that partition the Euclidean trajectory space. Further, because
 111 $x_{t+1} = x_t + \Delta x_t$, `baller2vec` is in fact a conditional generative model that assigns a probability
 112 to a sequence of trajectories given the initial position of the agent, i.e., $p(\Delta x_1, \Delta x_2, \dots, \Delta x_T | x_1)$.
 113 Using the chain rule, we decompose the joint probability of the trajectories as:

$$p(\Delta x_1, \Delta x_2, \dots, \Delta x_T) = p(\Delta x_1)p(\Delta x_2|\Delta x_1) \dots p(\Delta x_T|\Delta x_1, \Delta x_2, \dots, \Delta x_{T-1})$$

114 to reflect their temporal structure (we omit the conditional initial position term for brevity). Therefore,
 115 new trajectories can be generated from `baller2vec` with the following procedure (see Figure 3):

- 116 1. First, sample one of the n different mixture components using the mixture proportions
 117 output from the classifier f (i.e., `baller2vec`) conditioned on the agent’s current position,
 118 i.e., $i \sim \text{Categorical}(\pi_1, \pi_2, \dots, \pi_n)$ where $[\pi_1, \pi_2, \dots, \pi_n] = f(x_t)$.
- 119 2. Next, sample a trajectory from the uniform distribution associated with the sampled compo-
 120 nent, i.e., $\Delta x_t \sim \mathcal{U}(a_i, b_i)$.
- 121 3. Finally, add the sampled trajectory to the agent’s input position to generate the agent’s
 122 position at the start of the next time step, i.e., $x_{t+1} = x_t + \Delta x_t$.

123 Let $[\Delta x_{min}, \Delta x_{max}]$ be an interval on the real line such that any trajectory $\Delta x <$
 124 Δx_{min} or $\Delta x \geq \Delta x_{max}$ has zero density (i.e., such trajectories are humanly impossible).
 125 Let $\{[a_i, b_i]\}_{i=1}^n$ be a set of n intervals that par-
 126 tition the interval $[\Delta x_{min}, \Delta x_{max}]$ into n bins,
 127 i.e., $\cup_{i=1}^n [a_i, b_i] = [\Delta x_{min}, \Delta x_{max}]$ and $i \neq$
 128 $j \implies [a_i, b_i] \cap [a_j, b_j] = \emptyset$. Recall that the
 129 probability density function (PDF) for a uniform
 130 distribution with bounds $-\infty < a < b < \infty$ is:

$$p(\Delta x) = \begin{cases} \frac{1}{b-a} & \text{for } \Delta x \in [a, b] \\ 0 & \text{otherwise} \end{cases}$$

131 Letting $c_i = \frac{1}{b_i - a_i}$, the PDF for a mixture of
 132 uniforms with these bounds is thus:

$$p(\Delta x) = \sum_{i=1}^n \pi_i \mathcal{U}(\Delta x; a_i, b_i) = \sum_{i=1}^n \pi_i c_i \quad (1)$$

133 where $p(\Delta x)$ is the density assigned to Δx
 134 by the mixture, π_i is the mixture proportion
 135 for the mixture component indexed by i (i.e.,
 136 $0 \leq \pi_i \leq 1$ and $\sum \pi_i = 1$), and $\mathcal{U}(\Delta x; a_i, b_i)$
 137 is the density assigned to Δx by the uniform dis-
 138 tribution with bounds $-\infty < a_i < b_i < \infty$.
 139 Because the bounds of the uniform distribu-
 140 tions partition $[\Delta x_{min}, \Delta x_{max}]$, Equation (1)
 141 reduces to:

$$p(\Delta x) = \pi_{i'} c_{i'}$$

142 where $\Delta x \in [a_{i'}, b_{i'}]$ (because the other uni-
 143 form distributions will assign a density of zero to
 144 Δx). The likelihood for data D (with $|D| = N$)
 145 is then:

$$\mathcal{L}(D) = \prod_{j=1}^N p(\Delta x_j) = \prod_{j=1}^N \pi_{j,i'} c_{j,i'}$$




Figure 3: `baller2vec` can be viewed as a condi-
 tional generative model that assigns a probability
 to a sequence of trajectories given the initial po-
 sitions of the agents. Here, we show a graphi-
 cal model depiction of a `baller2vec` model that
 generates a sequence of one-dimensional trajec-
 tories for a single agent. Given the initial posi-
 tion of the agent (the circle containing x_1), one
 of n different uniform distributions (the square
 containing i_1) is sampled using the mixture pro-
 portions (π_i) output by `baller2vec` (f). The
 agent’s trajectory (the diamond containing Δx_1)
 is then sampled from the selected uniform distri-
 bution, which has bounds $-\infty < a_i < b_i < \infty$.
 At the start of the next time step, the agent’s
 position is $x_2 = x_1 + \Delta x_1$. Maximizing the
 likelihood of `baller2vec` as a classifier over the
 binned trajectories is thus equivalent to maxi-
 mizing its likelihood when assuming the trajec-
 tories are generated from a mixture of uniform
 distributions that partition the Euclidean trajec-
 tory space (see Section 2.2 for details).

146 where $\pi_{j,i'}$ is the mixture proportion assigned to the component with $\Delta x_j \in [a_{i'}, b_{i'})$ and $c_{j,i'}$ is the
 147 associated density. Taking the negative logarithm of the likelihood gives:

$$-\ln(\mathcal{L}(D)) = -\sum_{j=1}^N \ln(\pi_{j,i'}) - \sum_{j=1}^N \ln(c_{j,i'}) \quad (2)$$

148 Because the bounds are fixed, the second summation is a constant, and Equation (2) becomes:

$$-\ln(\mathcal{L}(D)) = -\sum_{j=1}^N \ln(\pi_{j,i'}) + C$$

149 where $C = -\sum_{j=1}^N \ln(c_{j,i'})$. Therefore, minimizing the loss of `baller2vec` as a classifier of
 150 binned trajectories is equivalent to minimizing the loss of the model when assuming the trajectories
 151 are generated from a mixture of uniform distributions as specified in Equation (1).

152 3 Model Architecture

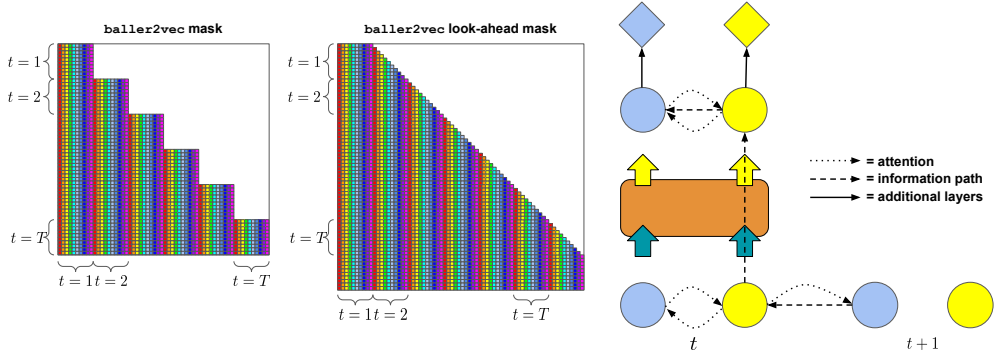


Figure 4: A naive strategy for learning to predict statistically dependent agent trajectories is to adapt the `baller2vec` self-attention mask so that `baller2vec` can “look ahead” at future positions of agents whose trajectories are generated *prior* to the agent being processed in the current time step. However, this look-ahead self-attention mask cannot be used with multi-layer Transformers because doing so necessitates “seeing the future”. For example, after the model attends to the blue agent’s position at time step $t + 1$ when processing the yellow agent at time step t , the yellow agent’s resultant feature vector contains information about the blue agent’s future position. As a result, when the model attends to the yellow agent while processing the blue agent at the next level, the model is seeing the future.

153 We motivate our `baller2vec++` architecture by first highlighting an issue that arises in `baller2vec`
 154 when trying to model agent trajectories using the chain rule. The `baller2vec` self-attention mask
 155 can be adapted so that `baller2vec` “looks ahead” at the future positions of agents whose trajectories
 156 are generated prior to the agent being processed in the current time step (Figure 4). However, this
 157 look-ahead self-attention mask can only be used with the final layer of the Transformer; otherwise,
 158 the model needs to see the future (Figure 4). As a result, `baller2vec` is severely limited in the
 159 conditional distribution functions it can learn.

160 `baller2vec++` (Figure 5) overcomes this limitation by: (1) augmenting the `baller2vec` input with
 161 two other sets of feature vectors and (2) using a specially designed self-attention mask. The three sets
 162 of feature vectors in `baller2vec++` take the following forms:

- 163 1. $z_{t,k} = g_z([e(p_k), x_{t,k}, y_{t,k}, h_{t,k}])$ (current location information)
- 164 2. $u_{t,k} = g_u([e(p_k), x_{t+1,k}, y_{t+1,k}, h_{t,k}, \Delta x_{t,k}, \Delta y_{t,k}])$ (“look-ahead” information)
- 165 3. $r_k = g_r([e(p_k), x_{1,k}, y_{1,k}, h_{1,k}])$ (initial location information)

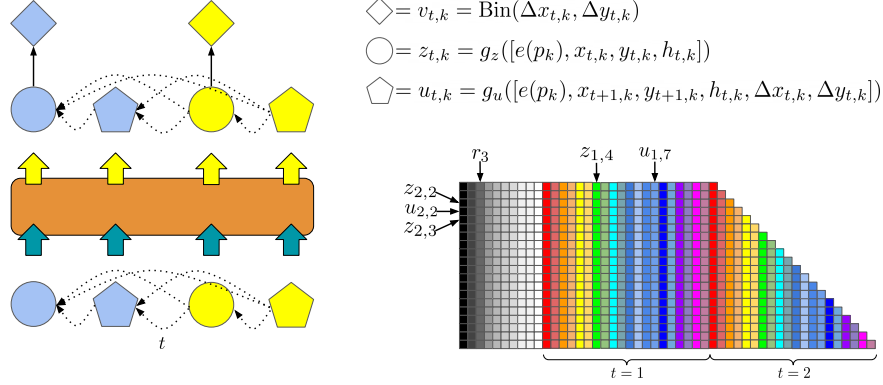


Figure 5: To learn statistically dependent agent trajectories, `baller2vec++` uses a specially designed self-attention mask to simultaneously process three different sets of features vectors in a single Transformer. The three sets of feature vectors consist of location feature vectors like those found in `baller2vec` ($z_{t,k}$), look-ahead trajectory feature vectors ($u_{t,k}$), and starting location feature vectors (r_k ; not shown). As can be seen in these partial depictions of `baller2vec++` and the `baller2vec++` self-attention mask, this design allows the model to integrate information about *concurrent* agent trajectories through *multiple* Transformer layers without seeing the future.

166 where g_z , g_u , and g_r are multilayer perceptrons (MLPs), e is an agent embedding layer, and $h_{t,k}$ is a
 167 vector of optional contextual features for agent p_k at time step t . $z_{t,k}$ is the same location feature
 168 vector used in `baller2vec` and contains information about a specific agent’s identity and the agent’s
 169 location at time step t . $u_{t,k}$ is a “look-ahead” trajectory feature vector that contains information about
 170 a specific agent’s identity, the agent’s location at the *next time step* $t + 1$, and the agent’s trajectory
 171 at time step t , i.e., $(x_{t+1,k} - x_{t,k}, y_{t+1,k} - y_{t,k})$. Lastly, r_k is a starting location feature vector that
 172 contains information about a specific agent’s identity and the agent’s location at time step $t = 1$.
 173 The r_k feature vectors are necessary so that `baller2vec++` can “see” the initial locations of *all the*
 174 *agents* when processing the agents agent-wise in the first time step.

175 These three sets of feature vectors are combined to form a $(K + 2TK) \times F$ matrix Z such that
 176 the first K rows consist of the K r_k feature vectors, and the remaining $2TK$ rows consist of the
 177 TK $z_{t,k}$ and TK $u_{t,k}$ feature vectors interleaved with one another, i.e., each $z_{t,k}$ is followed by its
 178 corresponding $u_{t,k}$ in the matrix. This matrix is passed into the Transformer along with the specially
 179 designed self-attention mask, which encodes the following dependencies (see Figure 5):

- 180 1. When processing r_{k_1} , `baller2vec++` is exclusively allowed to “look” at each r_{k_2} (i.e.,
 181 `baller2vec++` cannot look at any location or look-ahead feature vectors when processing
 182 r_{k_1}).
- 183 2. When processing z_{t_2,k_2} , `baller2vec++` is allowed to “look” at: (i) each r_{k_1} , (ii) any z_{t_1,k_1}
 184 where (a) $t_1 < t_2$ **or** (b) $t_1 = t_2$ **and** $k_1 \leq k_2$, and (iii) any u_{t_1,k_1} where (a) $t_1 < t_2$ **or** (b)
 185 $t_1 = t_2$ **and** $k_1 < k_2$.
- 186 3. When processing u_{t_2,k_2} , `baller2vec++` is allowed to “look” at: (i) each r_{k_1} , (ii) any z_{t_1,k_1}
 187 where (a) $t_1 < t_2$ **or** (b) $t_1 = t_2$ **and** $k_1 \leq k_2$, and (iii) any u_{t_1,k_1} where (a) $t_1 < t_2$ **or** (b)
 188 $t_1 = t_2$ **and** $k_1 \leq k_2$.

189 Each processed $z_{t,k}$ feature vector is then passed through a linear layer that is followed by a softmax,
 190 which gives a probability distribution over the trajectory bins for agent p_k at time step t . Similar to
 191 `baller2vec`, the loss for each sample is:

$$\mathcal{L} = \sum_{t=1}^T \sum_{k=1}^K -\ln(f(Z)_{t,2k-1}[v_{t,k}]) \quad (3)$$

192 where $f(Z)_{t,2k-1}[v_{t,k}]$ is the probability assigned to the trajectory bin $v_{t,k}$ (where $v_{t,k} =$
 193 $\text{Bin}(\Delta x_{t,k}, \Delta y_{t,k})$ is an integer from one to n^2) by f , i.e., Equation (3) is the negative log-likelihood
 194 (NLL) of the data according to the model.

195 Because any ordering of a chain rule decomposition of a joint probability produces the same value,
196 e.g.:

$$p(\Delta x_{t,1})p(\Delta x_{t,2}|\Delta x_{t,1})p(\Delta x_{t,3}|\Delta x_{t,1}\Delta x_{t,2}) = p(\Delta x_{t,3})p(\Delta x_{t,2}|\Delta x_{t,3})p(\Delta x_{t,1}|\Delta x_{t,3}\Delta x_{t,2})$$

197 like [11], we shuffled the order of the agents in each training sequence to encourage the model to
198 learn joint probabilities of the agent trajectories that are approximately permutation invariant with
199 respect to the ordering of the agents.

200 4 Experiments

201 We tested `baller2vec++` on two different datasets. To highlight the pathological behavior of models
202 that assume agent trajectories are statistically independent at each time step, we trained scaled down
203 versions of `baller2vec++` and `baller2vec` on a toy dataset consisting of simulated trajectories
204 for two perfectly coordinated agents. Additionally, to demonstrate the efficacy of `baller2vec++`
205 in real world settings, we trained `baller2vec++` and `baller2vec` on a dataset of trajectories for
206 professional basketball players.

207 4.1 Toy dataset

208 Each training sample was initialized with the agents starting at $(-1, 0)$ and $(1, 0)$ on a grid in random
209 order (i.e., the first agent could be placed to either the left or the right of the origin). At each time
210 step, one of nine actions (corresponding to the 3×3 grid surrounding the agent) was sampled from a
211 uniform distribution, and each of the agents was translated along this trajectory. This process was
212 repeated for 20 time steps (see Figure 6(a) for a sample).

213 4.2 Basketball dataset

214 We used the same National Basketball Association (NBA) dataset³ employed by Alcorn and Nguyen
215 [6], whom we paraphrase here:

216 The NBA dataset consists of trajectories from 631 games from the 2015-2016
217 season, which were split into 569/30/32 training/validation/test games, respectively.
218 During training, each sequence was sampled using the following procedure: (1)
219 randomly select a training game, (2) randomly select a starting time from the game,
220 (3) take the following four seconds of data and downsample it to 5 Hz from the
221 original 25 Hz, and then (4) randomly (with a probability of 0.5) rotate the court
222 180° . This sampling procedure gave us access to on the order of ~ 82 million
223 different (albeit overlapping) training sequences. For both the validation and test
224 sets, $\sim 1,000$ different, *non-overlapping* sequences were selected for evaluation by
225 dividing each game into $\lceil \frac{1,000}{N} \rceil$ non-overlapping chunks (where N is the number
226 of games), and using the starting four seconds from each chunk as the evaluation
227 sequence.

228 4.3 Model

229 Our `baller2vec++` and `baller2vec` models for the basketball dataset closely followed [6], and so
230 largely resemble the original Transformer architecture [12]. Specifically, the Transformer settings
231 were: $d_{\text{model}} = 512$ (the dimension of the input and output of each Transformer layer), eight attention
232 heads, $d_{\text{ff}} = 2048$ (the dimension of the inner feedforward layers), six layers, no dropout, and no
233 positional encoding. Each MLP (i.e., g_z , g_u , and g_r) had 128, 256, and 512 nodes in its three layers,
234 respectively, and a ReLU nonlinearity following each of the first two layers. The player embeddings
235 [13] had 20 dimensions, and $h_{t,k}$ was a binary variable indicating the side of the frontcourt for
236 player p_k (i.e., the direction of his team’s hoop) at time step t . Lastly, the $11 \text{ ft} \times 11 \text{ ft}$ 2D Euclidean
237 trajectory space was binned into $121 \text{ ft} \times 1 \text{ ft}$ squares.

238 We used the Adam optimizer [14] with an initial learning rate of 10^{-6} , $\beta_1 = 0.9$, $\beta_2 = 0.999$, and
239 $\epsilon = 10^{-9}$ to update the model parameters, of which there were ~ 19 million. The learning rate was

³<https://github.com/linouk23/NBA-Player-Movements>

240 reduced to 10^{-7} after 20 epochs of the validation loss not improving. Models were implemented in
 241 PyTorch and trained on a single NVIDIA GTX 1080 Ti GPU for ~ 650 epochs (seven days) where
 242 each epoch consisted of 20,000 training samples, and the validation set was used for early stopping.

243 For the toy dataset, we used scaled down versions of the basketball models with $d_{\text{model}} = 128$, four
 244 attention heads, $d_{\text{ff}} = 512$, and two layers in the Transformer. Additionally, each MLP had two layers
 245 with 64 and 128 nodes, respectively. The models were trained for 50 epochs of 500 samples per
 246 epoch (~ 10.5 minutes) using a single learning rate of 10^{-5} .

247 5 Results

248 5.1 baller2vec++ can effectively model coordinated agents in both simulated and real 249 settings

250 For the toy dataset, the training loss for `baller2vec` converged to $\sim 2.2 \approx -\ln(\frac{1}{9})$, i.e., the model
 251 was simply independently guessing the trajectories for both agents at every time step. In contrast,
 252 the training loss for `baller2vec++` converged to $\sim 1.1 \approx -\ln(\frac{1}{9}) \div 2$, which is the expected loss
 253 for a model that perfectly learns the deterministic relationship between the agents’ trajectories
 254 (because the prediction for the second agent will always contribute $-\ln(1.0) = 0$ to the loss).

255 When generating trajec-
 256 tories with `baller2vec`, the
 257 agents are completely unco-
 258 ordinated, with each agent
 259 following an independent
 260 random walk around the
 261 grid (Figure 6(b)). In con-
 262 trast, trajectories generated
 263 by `baller2vec++` display
 264 the same coordinated agent
 265 behavior as the training data
 266 (Figure 6(c)).




Figure 6: When trained on a dataset of perfectly coordinated agent tra-
 267 jectories (a), the trajectories generated by `baller2vec` are completely
 268 *uncoordinated* (b) while the trajectories generated by `baller2vec++`
 269 are perfectly coordinated (c). Animated versions can be found in the
 270 code repository.

271 For the basketball dataset, `baller2vec++` achieved
 272 an average NLL of 0.472
 273 on the test set, 8.9% better
 274 than the average NLL for
 275 `baller2vec` (0.518) (see Figure S1 for trajectories generated by `baller2vec++` and `baller2vec`).
 276 As was observed in [6], the trajectory bin distributions for `baller2vec` become much more certain
 277 after observing a portion of the sequence (Figure 7), which suggests `baller2vec` may be inferring
 278 some of the shared hidden variables (e.g., plays) influencing the players. If that hypothesis was true,
 279 the performance gap between `baller2vec++` and `baller2vec` should be largest at the beginning
 280 of the sequence (before any shared hidden variables can be inferred by `baller2vec`). Indeed, the
 281 average NLL for `baller2vec++` in the *first time step* of each test set sequence (1.567) is 16.1%
 282 better than the average NLL for `baller2vec` (1.869), while the average NLL for `baller2vec++` in
 283 the *last time step* of each test set sequence (0.420) is only 9.7% better than the average NLL for
 284 `baller2vec` (0.465) (see Figure 7).

282 5.2 baller2vec++ makes better predictions when conditioned on concurrent trajectory 283 information from other agents

284 Implicit in much of our discussion has been the intuition that providing a model with additional
 285 (relevant) information will improve its performance. To empirically test this conjecture, we compared
 286 the performance of `baller2vec++` when predicting the trajectory of a specific basketball player
 287 placed in the *first position* of the player order (i.e., when $k = 1$) vs. predicting the trajectory for that
 288 *same player* placed in the *last position* (i.e., when $k = 10$). Specifically, for each player in each test
 289 sequence, we calculated the NLL of the player’s trajectory in the first time step⁴ with the player in the

⁴Because, as previously discussed, the benefits of `baller2vec++` were most pronounced in the first time
 step.




Figure 7: **Left:** when modeling the trajectories of professional basketball players, the performance gap between `baller2vec++` and `baller2vec` is largest at the beginning of the sequence, before shared unobserved variables can be inferred by `baller2vec`. Each bar indicates a model’s average NLL over the entire test set for that particular time step. For full sequences, `baller2vec++` outperforms `baller2vec` by 8.9%. **Right:** the joint probability assigned to a sequence by `baller2vec++` is approximately permutation invariant with respect to the order of the agents. For each point, its x value indicates `baller2vec++`’s average NLL for a test set sequence using the *original* order of the agents in the sequence, while its y value indicates `baller2vec++`’s average NLL *for the same sequence* with the order of the agents *shuffled*. The shuffled average NLLs are highly correlated with their corresponding unshuffled average NLLs.

290 *first position* of the player order. Next, we moved the player to the *last position* of the player order,
 291 and then randomly shuffled the remaining nine players 10 times, calculating the NLL for the player
 292 in the last position each time. Finally, we calculated the average percent change in the last position
 293 NLLs relative to their corresponding first position NLLs. On average, moving a player from the first
 294 to the last position improved the NLL for the player’s trajectory by 14.6%.

295 **5.3 The joint probability assigned to a sequence by `baller2vec++` is approximately**
 296 **permutation invariant with respect to the order of the agents**

297 To determine whether or not `baller2vec++` respects the fact that any ordering of a chain rule
 298 decomposition of a joint probability produces the same value, we measured how much the average
 299 NLL for each test sequence in the basketball dataset varied when the order of the agents changed.
 300 Specifically, for each test set sequence, we shuffled the order of the agents 10 times. Then, for each
 301 permuted sequence, we calculated the percent error⁵ in the average NLL relative to the original,
 302 *unshuffled* sequence. Across all test sequences, the average percent error was only $\pm 1.5\%$. Further,
 303 as can be seen in Figure 7, the shuffled average NLLs are highly correlated with their corresponding
 304 unshuffled average NLLs (Pearson correlation coefficient = 0.997), i.e., the joint probability assigned
 305 to a sequence by `baller2vec++` is approximately permutation invariant with respect to the order of
 306 the agents.

307 **6 Conclusion and Future Work**

308 In this paper, we have shown how the commonly used independence assumption of many multi-agent
 309 spatiotemporal models can severely limit their ability to learn to emulate coordinated agents. By
 310 relaxing this independence assumption in `baller2vec`, `baller2vec++` was able to more accurately
 311 model the trajectories of professional basketball players. Models for other multi-agent spatiotemporal
 312 environments, such as pedestrian traffic (see [15] for a survey) and vehicle traffic (e.g., [16–19]), may
 313 also benefit from the look-ahead approach used by `baller2vec++`. Additionally, the interleaved
 314 input design of `baller2vec++` could be useful when modeling other systems involving many entities
 315 interacting through time, such as social media platforms (e.g., [20, 21]). However, confronting the
 316 quadratic complexity of the Transformer attention mechanism as the number of entities grows large
 317 in these datasets is an open problem, but recent work in sparse Transformers (e.g., [22–27]) shows
 318 encouraging progress.

⁵See: https://en.wikipedia.org/wiki/Approximation_error#Formal_Definition.

References

- 319
- 320 [1] Panna Felsen, Patrick Lucey, and Sujoy Ganguly. Where will they go? predicting fine-grained
321 adversarial multi-agent motion using conditional variational autoencoders. In *Proceedings of*
322 *the European Conference on Computer Vision (ECCV)*, pages 732–747, 2018.
- 323 [2] Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi. Social gan:
324 Socially acceptable trajectories with generative adversarial networks. In *Proceedings of the*
325 *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2255–2264, 2018.
- 326 [3] Amir Sadeghian, Vineet Kosaraju, Ali Sadeghian, Noriaki Hirose, Hamid Rezaatofghi, and
327 Silvio Savarese. Sophie: An attentive gan for predicting paths compliant to social and physical
328 constraints. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*
329 *Recognition*, pages 1349–1358, 2019.
- 330 [4] Raymond A Yeh, Alexander G Schwing, Jonathan Huang, and Kevin Murphy. Diverse genera-
331 tion for multi-agent sports games. In *Proceedings of the IEEE Conference on Computer Vision*
332 *and Pattern Recognition*, pages 4610–4619, 2019.
- 333 [5] Cunjun Yu, Xiao Ma, Jiawei Ren, Haiyu Zhao, and Shuai Yi. Spatio-temporal graph transformer
334 networks for pedestrian trajectory prediction. In *Proceedings of the European Conference on*
335 *Computer Vision (ECCV)*, August 2020.
- 336 [6] Michael A. Alcorn and Anh Nguyen. baller2vec: A multi-entity transformer for multi-agent
337 spatiotemporal modeling. *arXiv preprint arXiv:2102.03291*, 2021.
- 338 [7] Eric Zhan, Stephan Zheng, Yisong Yue, Long Sha, and Patrick Lucey. Generating multi-agent
339 trajectories using programmatic weak supervision. In *International Conference on Learning*
340 *Representations*, 2019. URL <https://openreview.net/forum?id=rkxw-hAcFQ>.
- 341 [8] Stephan Zheng, Yisong Yue, and Jennifer Hobbs. Generating long-term trajectories using deep
342 hierarchical networks. *Advances in Neural Information Processing Systems*, 29:1543–1551,
343 2016.
- 344 [9] Vineet Kosaraju, Amir Sadeghian, Roberto Martín-Martín, Ian Reid, Hamid Rezaatofghi, and
345 Silvio Savarese. Social-bigat: Multimodal trajectory forecasting using bicycle-gan and graph
346 attention networks. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox,
347 and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32.
348 Curran Associates, Inc., 2019. URL [https://proceedings.neurips.cc/paper/2019/](https://proceedings.neurips.cc/paper/2019/file/d09bf41544a3365a46c9077ebb5e35c3-Paper.pdf)
349 [file/d09bf41544a3365a46c9077ebb5e35c3-Paper.pdf](https://proceedings.neurips.cc/paper/2019/file/d09bf41544a3365a46c9077ebb5e35c3-Paper.pdf).
- 350 [10] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua
351 Bengio. Graph Attention Networks. *International Conference on Learning Representations*,
352 2018. URL <https://openreview.net/forum?id=rJXMpikCZ>. accepted as poster.
- 353 [11] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V
354 Le. Xlnet: Generalized autoregressive pretraining for language understanding. *arXiv preprint*
355 *arXiv:1906.08237*, 2019.
- 356 [12] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,
357 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Informa-*
358 *tion Processing Systems*, pages 5998–6008, 2017.
- 359 [13] Michael A Alcorn. (batter|pitcher)2vec: Statistic-free talent modeling with neural player
360 embeddings. In *MIT Sloan Sports Analytics Conference*, 2018.
- 361 [14] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Interna-*
362 *tional Conference on Learning Representations*, 2015.
- 363 [15] Andrey Rudenko, Luigi Palmieri, Michael Herman, Kris M Kitani, Darius M Gavrilă, and Kai O
364 Arras. Human motion trajectory prediction: A survey. *The International Journal of Robotics*
365 *Research*, 39(8):895–935, 2020.

- 366 [16] Nachiket Deo and Mohan M Trivedi. Convolutional social pooling for vehicle trajectory
367 prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*
368 *Workshops*, pages 1468–1476, 2018.
- 369 [17] Ming-Fang Chang, John Lambert, Patsorn Sangkloy, Jagjeet Singh, Slawomir Bak, Andrew
370 Hartnett, De Wang, Peter Carr, Simon Lucey, Deva Ramanan, and James Hays. Argoverse:
371 3d tracking and forecasting with rich maps. In *Proceedings of the IEEE/CVF Conference on*
372 *Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- 373 [18] Tianyang Zhao, Yifei Xu, Mathew Monfort, Wongun Choi, Chris Baker, Yibiao Zhao, Yizhou
374 Wang, and Ying Nian Wu. Multi-agent tensor fusion for contextual trajectory prediction. In
375 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*,
376 June 2019.
- 377 [19] Rohan Chandra, Uttaran Bhattacharya, Aniket Bera, and Dinesh Manocha. Traphic: Trajectory
378 prediction in dense and heterogeneous traffic using weighted interactions. In *Proceedings of the*
379 *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- 380 [20] Srijan Kumar, Xikun Zhang, and Jure Leskovec. Predicting dynamic embedding trajectory
381 in temporal interaction networks. In *Proceedings of the 25th ACM SIGKDD International*
382 *Conference on Knowledge Discovery & Data Mining*, pages 1269–1278, 2019.
- 383 [21] Emanuele Rossi, Ben Chamberlain, Fabrizio Frasca, Davide Eynard, Federico Monti, and
384 Michael Bronstein. Temporal graph networks for deep learning on dynamic graphs. In *ICML*
385 *2020 Workshop on Graph Representation Learning*, 2020.
- 386 [22] Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. Generating long sequences with
387 sparse transformers. *arXiv preprint arXiv:1904.10509*, 2019.
- 388 [23] Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti,
389 Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed. Big
390 bird: Transformers for longer sequences. In H. Larochelle, M. Ranzato, R. Hadsell, M. F.
391 Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33,
392 pages 17283–17297. Curran Associates, Inc., 2020. URL [https://proceedings.neurips.
393 cc/paper/2020/file/c8512d142a2d849725f31a9a7a361ab9-Paper.pdf](https://proceedings.neurips.cc/paper/2020/file/c8512d142a2d849725f31a9a7a361ab9-Paper.pdf).
- 394 [24] Iz Beltagy, Matthew E Peters, and Arman Cohan. Longformer: The long-document transformer.
395 *arXiv preprint arXiv:2004.05150*, 2020.
- 396 [25] Nikita Kitaev, Lukasz Kaiser, and Anselm Levskaya. Reformer: The efficient transformer. In
397 *International Conference on Learning Representations*, 2020. URL [https://openreview.
398 net/forum?id=rkgNkkHtvB](https://openreview.net/forum?id=rkgNkkHtvB).
- 399 [26] Jonathan Ho, Nal Kalchbrenner, Dirk Weissenborn, and Tim Salimans. Axial attention in
400 multidimensional transformers. *arXiv preprint arXiv:1912.12180*, 2019.
- 401 [27] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for
402 video understanding? *arXiv preprint arXiv:2102.05095*, 2021.

403 **Checklist**

- 404 1. For all authors...
- 405 (a) Do the main claims made in the abstract and introduction accurately reflect the paper's
406 contributions and scope? [Yes]
- 407 (b) Did you describe the limitations of your work? [Yes] See Section 6.
- 408 (c) Did you discuss any potential negative societal impacts of your work? [No] Our work
409 does not introduce new ethical challenges.
- 410 (d) Have you read the ethics review guidelines and ensured that your paper conforms to
411 them? [Yes]
- 412 2. If you are including theoretical results...
- 413 (a) Did you state the full set of assumptions of all theoretical results? [Yes]
- 414 (b) Did you include complete proofs of all theoretical results? [Yes]
- 415 3. If you ran experiments...
- 416 (a) Did you include the code, data, and instructions needed to reproduce the main experi-
417 mental results (either in the supplemental material or as a URL)? [Yes]
- 418 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they
419 were chosen)? [Yes]
- 420 (c) Did you report error bars (e.g., with respect to the random seed after running experi-
421 ments multiple times)? [N/A]
- 422 (d) Did you include the total amount of compute and the type of resources used (e.g., type
423 of GPUs, internal cluster, or cloud provider)? [Yes]
- 424 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 425 (a) If your work uses existing assets, did you cite the creators? [Yes]
- 426 (b) Did you mention the license of the assets? [No] We link directly to the dataset.
- 427 (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
- 428 (d) Did you discuss whether and how consent was obtained from people whose data you're
429 using/curating? [N/A]
- 430 (e) Did you discuss whether the data you are using/curating contains personally identifiable
431 information or offensive content? [N/A]
- 432 5. If you used crowdsourcing or conducted research with human subjects...
- 433 (a) Did you include the full text of instructions given to participants and screenshots, if
434 applicable? [N/A]
- 435 (b) Did you describe any potential participant risks, with links to Institutional Review
436 Board (IRB) approvals, if applicable? [N/A]
- 437 (c) Did you include the estimated hourly wage paid to participants and the total amount
438 spent on participant compensation? [N/A]