Authors' reply to the Discussion of 'Flexible marked spatio-temporal point process with applications to event sequences from association football' by Narayanan, Kosmidis and Dellaportas

September 21, 2023

1 Introduction

We are grateful to all discussants for the insightful comments, thoughts, pointers and proposals for extending our work. We identified that the discussion contributions evolve around the themes: 1) Estimation and scalability, 2) Evaluation of predictive performance and overfitting, 3) Choice of field zones, 4) Inclusion of extra covariate information, 5) Game simulation and in-game forecasts, 6) Model outputs, 7) Data considerations, and 8) Model extensions. In what follows, we structure our reply according to those themes directly referring to the relevant contributions.

2 Estimation and scalability

Karlis (2023) and Egidi (2023) note that the current Hamiltonian Monte Carlo implementation can be slow and would benefit from improvement. We acknowledge these observations and note that this has been dealt with in the recent work by Panos et al. (2023) that develops a variational inference framework to provide a highly scalable procedure for training the models we introduced. Panos et al. (2023) also extend the model to allow for time-varying abilities within their proposed variational inference framework and report computational times of a few hours for a whole season's worth of touch-ball data.

Stival and Schiavon (2023) suggested using sparsity-inducing priors as an alternative to our work's association rule learning method to deal with model complexity. That is an excellent suggestion that we plan to pursue as part of future work. The main challenge we faced in our limited attempts with sparsity-inducing priors within the current vanilla posterior sampling framework is again the dimension of the parameter space. Nevertheless, we believe such prior structures can prove helpful alongside the variational inference framework of Panos et al. (2023).

3 Evaluation of predictive performance and overfitting

Egidi (2023) enquires about posterior predictive checks to assess model accuracy and the amount of overfitting in the model predictions.

In our work, we evaluate the predictive performance of the models using the log-pointwise predictive density on test data; see Section 6.2 of the main text. Of course, if prediction is the aim of the modelling exercise, we recognise that it is helpful to evaluate the models' predictive performance using additional evaluation criteria that compare predictions against observables more explicitly. Examples include out-of-sample root mean square error (RMSE) for the ability to predict the times of future events, and the wealth of classification performance measures for the ability to predict future marks (see, for example, Sokolova and Lapalme, 2009, for a systematic analysis of such). Within their variational framework, Panos et al. (2023) evaluated the model's performance using such criteria and found that it offers highly competitive computational and predictive performance against other state-of-the-art methods, which typically involve high-dimensional structures through semi- or non-parametric components.

We did not observe any severe overfitting of the models. The order of the models in terms of increasing *out-of-sample* log-pointwise predictive density in Table 7 is similar to the one obtained by naively evaluating the *in-sample* log-pointwise predictive density. The latter is an overestimate of the expected log-pointwise predictive density with respect to future data. A notable difference is with the MβA model (matrix β with team abilities), which has the largest number of parameters, and the largest *in-sample* log-pointwise predictive density. Despite being close to the best and second-best models, the MβA model is the third best in the *out-of-sample* evaluations in Table 7. That may indicate a small degree of overfitting for that model, most probably because, in the training data, each team plays just one game at their home and one at an away venue.

4 Choice of field zones

Karlis (2023), Egidi (2023) and Smith (2023) comment on the choice of partitioning the field into three zones of known area and enquire how that choice influences the model estimates.

The event triggering parameters of the Mβ parameterisation of the model (see expression (12) in the main text) depend on specifying a partition. That dependence enables us to readily infer the importance of different zones for particular actions, substantially enhancing interpretability (see Section 6.1-6.4 for such interpretations). For example, we can compute the chance of completing a successful pass or attempting a shot on goal in a particular zone. The choice of zones and their areas represent an idealisation of our understanding of the game, where the playing strategies have a natural dependence on whether the ball is in the defensive, midfield or attacking third of the field.

The recommendation of Smith (2023) for a more data-driven approach to determine the partitions is fruitful and an exciting area for future work. For example, we can consider a team-dependent, continuous spatial process that respects the field boundaries (see, for example, Solin and Kok, 2019) for $h(z_i \mid t_i, \mathcal{F}_{t_{i-1}}; \boldsymbol{\eta})$ in expression (7) of the main text, and threshold the field adaptively into a fixed number of zones. This way, zones will have an adaptive area that depends on how each team realises its strategy.

5 Inclusion of extra covariate information

Karlis (2023), Smith (2023) and Yurko and Nugent (2023) suggested including covariate information, such as in-game characteristics, off-the-ball player positions and player qualities to improve the model performance.

Our current case study uses only the team information as covariates to demonstrate the modelling framework. Nevertheless, the modelling framework readily allows for including other covariates to drive the cross-excitation of the marks; see Section 4.3 of the main text. Including covariates about the game's current state, such as the current score, number of cards, etc., may be particularly beneficial in predictive ability and is the topic of ongoing investigations. The cross-excitation of the marks can also incorporate information about the positions and qualities of players if that information is available.

Other parameters, such as the background process parameters, excitation factors and decay rates, can also be appropriately linked with covariate information through regression structures. The inferential or predictive benefits of including such regression structures should be weighed against the increased model complexity.

6 Game simulation and in-game forecasts

Karlis (2023) and Smith (2023) enquire about the model's ability to simulate whole games. The experiments in the paper only dealt with 30-second simulations.

The simulation framework of Section 7 of the manuscript and its implementation through the codebase we provide allows for the simulation of an arbitrary period or even the whole game using the in-game forecasts. We did run a limited number of simulations for entire games, and the forecasted scorelines were in the plausible spectrum. Still, more work is needed to study the usefulness of the proposed model for such applications.

7 Model outputs

Egidi (2023) notes that passing ability appears to be a markedly discriminant predictor of team rankings and wonders whether model-free passing ability statistics could be good for generating rankings. The model is not specified with the direct intent to predict team rankings. Our rankings are a bonus, but perhaps unsurprising, output from using a highly interpretable parameterisation. Further work is required to study the utility of the model parameters for generating team rankings using data from multiple seasons and leagues.

Yurko and Nugent (2023) asked if we explored using our approach to provide a real-time value for player decision-making. That is an exciting area of application, which we did not consider and for which our approach is well suited. Including player-level covariate information is necessary in that direction, and we intend to explore it in future work.

8 Data considerations

Egidi (2023) enquires whether a model trained on only two games for each of the 20 teams is stable enough. Our decision to use only the first two games for each of the 20 teams resulted from our attempt to demonstrate the wealth of insights that can be generated from our proposal using a limited amount of information, also accounting for the computational limitations we have been facing when fitting the models. A potential stability assessment could come by fitting the models over different sets of games. Note that the variational inference framework in Panos et al. (2023) overcomes the computational limitations and can fit the MβA model with time-varying abilities in a whole season's worth of touch-ball events in a few hours.

Smith (2023) suggests that the observations that events in football are more regular than Poisson may be because the shortest inter-event times are missing. That is true in the data analysed. Certain kinds of events in football, such as off-the-ball events like player runs, are not recorded. We found that the gamma model was adequate for modelling the inter-event times for the available data. Still, the truncated Poisson, as suggested by Smith (2023), is a valid alternative to compare with.

9 Model extensions

Mateu (2023) presents some helpful model extensions we have yet to consider and plan to investigate, including i) using periodicity on background rates and ii) using Hawkes models for the ball's trajectories. Naturally, and as identified by the discussant, the former extension is more direct through the conditional intensity function than through the decomposition of a multivariate distribution function in (4). A remedy is to employ similar specifications as in Zhuang and Mateu (2019) and add periodic effects on the definition background mark probabilities in (5). Extensions in direction ii) are directly possible through the appropriate specification of $h(z_i \mid t_i, \mathcal{F}_{t_{i-1}}; \boldsymbol{\eta})$ in expression (7) of the main text. See Section 4 for relevant discussion.

Another direction for extension mentioned by Mateu (2023) is accommodating negative interactions between events. Indeed, our development here cannot capture inhibition behaviour (i.e. having the occurrence of an event decrease the likelihood of another event to occur). We define a multivariate process on composite event types, where each event type is tracked for both the home and the away team (see Table 3 in the main text) to capture excitations from events both within each team and between teams. Capturing inhibition through our model specification is an exciting direction and valuable for the diverse applications of those models, for which recent developments such as Costa et al. (2020) and Bonnet et al. (2021) can be helpful.

Mateu (2023) also enquires how the probability of an event coming from the background or being triggered by another event can be computed. That is possible by calculating the posterior conditional branching structure probabilities in expression (17), which we use for deriving event genealogies in Section 6.8 of the main text.

References

Bonnet, A., M. Martinez Herrera, and M. Sangnier (2021). Maximum likelihood estimation for hawkes processes with self-excitation or inhibition. *Statistics & Probability Letters 179*, 109214.

Costa, M., C. Graham, L. Marsalle, and V. C. Tran (2020). Renewal in hawkes processes with self-excitation and inhibition. *Advances in Applied Probability 52*(3), 879–915.

Egidi, L. (2023). Seconder of the vote of thanks to Narayanan, Kosmidis and Dellaportas and contribution to the Discussion of "Flexible marked spatio-temporal point processes with applications to event sequences from association football". *Journal of the Royal Statistical Society Series C: Applied Statistics*.

Karlis, D. (2023). Proposer of the vote of thanks to Narayanan, Kosmidis and Dellaportas and contribution to the Discussion of "Flexible marked spatio-temporal point processes with applications to event sequences from association football". *Journal of the Royal Statistical Society Series C: Applied Statistics*.

Mateu, J. (2023). Jorge Mateu's contribution to the Discussion of 'Flexible marked spatio-temporal point processes with applications to event sequences from association football' by Narayanan, Kosmidis and Dellaportas. *Journal of the Royal Statistical Society Series C: Applied Statistics*.

Panos, A., I. Kosmidis, and P. Dellaportas (2023). Scalable marked point processes for exchangeable and non-exchangeable event sequences. In F. Ruiz, J. Dy, and J.-W. van de Meent (Eds.), *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, Volume 206 of *Proceedings of Machine Learning Research*, pp. 236–252. PMLR.

Smith, P. A. (2023). Paul Smith's contribution to the Discussion of "Flexible marked spatio-temporal point processes with applications to event sequences from association football" by Narayanan, Kosmidis and Dellaportas. *Journal of the Royal Statistical Society Series C: Applied Statistics*.

Sokolova, M. and G. Lapalme (2009). A systematic analysis of performance measures for classification tasks. *Information Processing & Management 45*(4), 427–437.

Solin, A. and M. Kok (2019). Know your boundaries: Constraining gaussian processes by variational harmonic features. In K. Chaudhuri and M. Sugiyama (Eds.), *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, Volume 89 of *Proceedings of Machine Learning Research*, pp. 2193–2202. PMLR.

Stival, M. and L. Schiavon (2023). Mattia Stival and Lorenzo Schiavon's contribution to the Discussion of "Flexible marked spatio-temporal point processes with applications to event sequences from association football" by Narayanan, Kosmidis and Dellaportas. *Journal of the Royal Statistical Society Series C: Applied Statistics*.

Yurko, R. and R. Nugent (2023). Ron Yurko and Rebecca Nugent's contribution to the Discussion of "Flexible marked spatio-temporal point processes with applications to event sequences from association football" by Narayanan, Kosmidis and Dellaportas. *Journal of the Royal Statistical Society Series C: Applied Statistics*.

Zhuang, J. and J. Mateu (2019). A semiparametric spatiotemporal hawkes-type point process model with periodic background for crime data. *Journal of the Royal Statistical Society Series A: Statistics in Society 182*(3), 919–942.