Proceedings of the 2021 Winter Simulation Conference
S. Kim, B. Feng, K. Smith, S. Masoud, Z. Zheng, C. Szabo, and M. Loper, eds.

MODELLING THE MENTEE-MENTOR POPULATION DYNAMICS: CONTINUOUS AND DISCRETE APPROACHES

Samuel Schaffel
François-Alex Bourque

Centre for Operational Research and Analysis
Development Research and Defence Canada
60 Moodie Drive
Ottawa, ON K1A 0K2, Canada

Slawomir Wesolkowski

Royal Canadian Air Force
Department of National Defence
101 Colonel By Drive
Ottawa, ON K1A 0K2, Canada

ABSTRACT

In order to be effective, the military workforce needs to continually produce experienced, well-trained personnel. Personnel go through different phases: recruitment (production), experience building (absorption) and retention (avoiding attrition). Absorption in military workforce modelling is often limited to maintaining the health of an occupation. The mentoring process, the conveying of experience to junior staff by senior staff, is rarely modelled to realistically represent the absorption of new personnel under varied occupational health conditions, including high ratios of mentees to mentors. To explore this dynamic, a two-state model based on the well-known predator-prey model is introduced where mentees upgrade to mentors, and the training capacity is modulated by the number of mentors. A continuous deterministic model is compared to a discrete stochastic simulation. The models consider a hypothetical military occupation in which the mentor-mentee relationship is important for career progression.

1 INTRODUCTION

People are the most important asset of the military enterprise. Military personnel play a critical role since, for example, without pilots aircraft will not fly and without sailors ships will not sail. To maintain military capabilities at a desired level, countries must make significant investments in their people: they must recruit the best (recruitment), train them (absorption), employ them in the most effective way (employment) and maintain their readiness (sustainment). Since personnel costs are usually the largest part of the defence budget, careful workforce planning has to be undertaken to ensure the military size and readiness correspond to a country's Level of Ambition (LoA), namely the number and type of operations it is able to carry out concurrently at any particular time.

To meet the country's LoA, the composition of the military workforce needs to produce experienced, well-trained personnel, which requires significant planning effort. These planning processes are complex because military workforces are comprised of hundreds of occupations requiring different degrees of specialization to support all the capabilities the military need to maintain, and because of the unique closed nature and strict rank hierarchy characteristic of the military workforce. Since the military fill higher ranks exclusively via internal promotions, in some specialized military trades, a new recruit will not be deemed fully employable until completing close to a decade of training. For this training to be successful, occupations can require mentorship of junior personnel by senior personnel. Thus, the level of experience (and not necessarily rank) of each soldier plays a critical part in how effectively military assets and infrastructure will be used and how effectively junior personnel will be trained. Therefore, unlike in most industries, military personnel production is constrained from both the supply and demand sides where a shortfall and a surplus of skilled staff can be costly and inefficient. This is a significant absorption issue.

978-1-6654-3311-2/21/$31.00 ©2021 Crown

Bastian and Hall provide an excellent overview of methods applied to military workforce planning and modelling (Bastian and Hall 2020). Foremost in the mathematical formulation of population dynamics are differential equations (Boileau 2012; Vincent and Okazawa 2019) and Markov decision models (Zais and Zhang 2015; Diener 2018; Suvorova, Novak, Moran, and Caelli 2019). Another approach is system dynamics (Forrester 1965) which utilizes continuous stocks and flows as well as feedback loops to examine how organizational structure, policies and decisions interact. In particular, a system dynamics approach was implemented to model mentee-mentor dynamics for pilots (Séguin 2015). Statistical methods have also been used to examine the composition of workforces and to project occupational requirements into the near future (Bryce and Henderson 2020; Okazawa 2020). Discrete-event simulations of the personnel flow within various armed forces occupations is, however, the most common method (Novak, Tracey, Nguyen, Johnstone, Le, and Creighton 2015; Henderson and Bryce 2019) for tracking the progress of individuals through a military training system.

With the notable exception of (Séguin 2015), most military models of workforce populations reviewed do not model the mentee-mentor dynamics beyond the usual flow upwards in ranks due to promotion. In addition, many treat the population variables as continuous including (Séguin 2015) in order to generate models that are computationally efficient and easier to use. However, workforce population models inherently deal with discrete entities, as one cannot have a fraction of a person hold a position. To improve our understanding of personnel absorption, we introduce population dynamics inspired from the well-known predator-prey model (Swift 2002). We present both continuous and discrete versions of the mentee-mentor population dynamic as well as an analysis of their differences.

The remainder of this paper is organized as follows. In Section 2, the problem is formulated and the two proposed models of mentee-mentor population dynamics are introduced. Section 3 compares the models for a benchmark scenario, while Section 4 pinpoints the dynamic effects at play. Next, Section 5 details which factors affect the difference between the two models. Finally, Section 6 concludes the paper.

2 MODELLING

When simulating the dynamics of a military workforce, the underlying assumption is often that the absorption of mentees is *not* affected by the number of mentors in the system (Boileau 2012; Vincent and Okazawa 2019; Zais and Zhang 2015; Diener 2018; Suvorova, Novak, Moran, and Caelli 2019). Therefore, the dynamics under the regime where the mentee-mentor ratio is high is usually omitted. However, while it may be the case that the desired healthy steady-state populations have a mentee-mentor ratio that does not limit the absorption of mentees, the system may still enter a high mentee-mentor ratio during the transition between two steady-states.

To further explore the behaviour of the system during this transient period, this paper explicitly models the mentee-mentor population dynamics by introducing a two-state model where mentees upgrade to mentors after a certain training period. The rate at which the mentees are trained depends on the number of mentors and their ability to mentor. This model may be applied to a variety of workplace settings; however, we focus specifically on the career progression of military occupations with on-the-job training such as pilot or aircraft technician occupations. The purpose is to track the expected populations of mentees, mentors and their progression while achieving the goal of having a healthy workforce in steady-state. In this state, the populations of mentees and mentors will be at their target sizes and the flow of mentees (i.e., new recruits) entering the workforce will balance out the flow of mentees upgrading to mentors and the flow of mentors leaving due to attrition.

In this paper, the same non-linear mentee-mentor dynamics are incorporated into: (1) a deterministic continuous model in Section 2.1; and (2) a discrete stochastic model in Section 2.2. The two models are then compared to determine how their outputs differ and the set of conditions under which the continuous model produces similar results to the discrete model.

2.1 Deterministic Continuous Model

The continuous mentee-mentor population dynamics are implemented using two non-linear ordinary differential equations: (1) dictates the population dynamics of the mentees (x), while (2) that of the mentors (y).

$$\dot{x} = a - b\min(x, ry) \tag{1}$$

$$\dot{y} = b\min(x, ry) - cy \tag{2}$$

As seen, in (1), the rate of change in population for the mentees ($\dot{x}$) is the intake rate (a) minus the absorption rate [$b\min(x, ry)$], and in (2) the rate of change for mentors ($\dot{y}$) is the absorption rate minus the attrition rate (cy). When the number of mentees is less than the training capacity of the mentors (ry), where r is the latent absorption capacity, then the absorption rate is only dependent on the number of mentees and the average time it takes for them to upgrade ($1/b$ time units). This is called the unsaturated regime because the training capacity of the mentors is not saturated. However, if x exceeds ry, then the system changes to the saturated regime since the mentors no longer have the capacity to train mentees at the same rate. As a result, mentors will have to spread their training capacity equally among all mentees, resulting in a longer training time for each individual mentee (since they are receiving less attention from the mentors). An alternative interpretation for this behavior is that mentees who arrive when the system is in the saturated regime have to wait in a queue until a spot opens up. Either way, the absorption rate no longer depends on the number of mentees, but rather on the number of mentors. This saturation effect in the training capacity of a mentor is the principal difference with the original predator-prey model where predators are assumed to have no satiety limits (Swift 2002).

2.2 Stochastic Discrete Model

The continuous model may be insufficient at capturing the workplace population dynamics. First, because it is a deterministic model, stochastic effects are absent. For a large population in steady state, this should not be noticeable; however, for smaller populations, since the fluctuations might represent a larger proportion of the total, their effects will be more pronounced. For example, a small population might, in some cases, be reduced to zero due to an increase in attrition. Once the population of mentors is eliminated, mentees can no longer be trained and the experience building system will grind to a halt. Second, using the continuous model may be inappropriate since continuously applying changes to the populations may have a different effect than discretely applying them. Thus, the need to explore a discrete stochastic model.

In this paper, a discrete-event Monte-Carlo simulation is implemented using the Gillespie algorithm (Mehta 2019). Specifically, there are three possible events:

1. A mentee can enter the system, resulting in an increase of the mentee population by one;
2. A mentee can upgrade to a mentor, resulting in the mentee population decreasing and the mentor population increasing by one; and
3. A mentor can leave the system, resulting in a decrease in the mentor population by one.

These events occur as Poisson processes with rates a, $b\min(x, ry)$, and cy respectively.

The Gillespie algorithm allows the option for the next-event technique; in this approach time jumps from event to event. This is possible under the assumption that all stochastic processes are Poisson. If each individual event occurs as a Poisson process with rate r_i, then the superposition of these is a Poisson process with rate equal to $r_{tot}=\Sigma\ r_i$, and it is from this process that the time between events is drawn. Furthermore, at each step of the Poisson process, the probability that event i occurs is equal to r_i/r_{tot} (Mehta 2019). Therefore, when a change in state happens, a random number generator can choose which specific event occurred. Then, the populations are updated and the simulation moves on to the next instance. The

Gillespie method makes the assumption that all events occur as independent Poisson processes. There are variations of the Gillespie algorithm which do not require this assumption (Masuda and Rocha 2020).

3 CONTINUOUS AND DISCRETE MODEL COMPARISON

As stated in Section 2, the continuous model may be insufficient at capturing the population dynamic. As such, the significance of the discrete effects can be quantified by comparing the discrete model with the continuous model. To benchmark the models, this study uses a realistic scenario. Specifically, the mentee and mentor populations transition from an initial steady-state (x_0, y_0) to a new target steady-state (x_{target}, y_{target}) where populations are doubled. The parameters are chosen such that they emulate numbers representative of an actual set of military trades. Note that the attrition constant c is treated as given, as it is considered a property of the system dynamics, while the intake a and absorption b constants are calculated to achieve the desired steady-state. Both the initial and target steady-states are also assumed to be healthy, i.e., to exhibit an unsaturated training regime. With that in mind, the relationship between system parameters is found by solving: $\dot{x} = \dot{y} = 0$. This gives the relation,

$$a = bx_{target} = cy_{target} \tag{3}$$

With this in mind, the parameters are set to:

$$\begin{aligned}
(x_0, y_0) &= (25, 100) \\
(x_{target}, y_{target}) &= (50, 200) \\
c &= 0.05 \\
r &= \frac{1}{3} \\
a &= cy_{target} = 10 \\
b &= \frac{a}{x_{target}} = 0.2
\end{aligned}$$

Figure 1 shows the results of the continuous model superimposed over the discrete model statistics (Figure 1a). As well, for illustrative purposes, the first three iterations of the discrete model (Figure 1b) are reported to show some possible system excursions. Figure 1a reports statistics of the discrete model simulated over 10,000 iterations. Specially, the mean (dotted lines) and median (dashed lines) of both the mentee and mentor populations over time, as well as the 25th-75th quantile envelope are plotted. While both the continuous and discrete models achieve the desired steady state in Figure 1a, there are significant differences between the two during the transient period between the initial and target populations. For the continuous model, the mentee and mentor populations evolve deterministically over time. The solid purple and green curves show the mentee and mentor population respectively, while the solid yellow capacity curve is the boundary between the two training regimes (ry). When the mentee curve lies above the capacity curve, the system is under the saturated regime, and when the mentee curve is below the capacity curve, the system is in the unsaturated regime. It is observed that the system quickly enters the saturated regime as there are not enough mentors to keep up with the intake of mentees. The mentee population then continues to increase until $y = a/br = 150$ (this value is found by solving (1) at $\dot{x} = 0$ in the saturated regime); at that point, the mentee population decreases, crosses back into the unsaturated regime, and settles to the steady state.

In comparison to the continuous model, the mean and median mentor curves take longer to reach the steady-state value. The difference between the two models is more pronounced when looking at the mentee curves where the mean of the discrete model overshoots the continuous model by a significant amount and

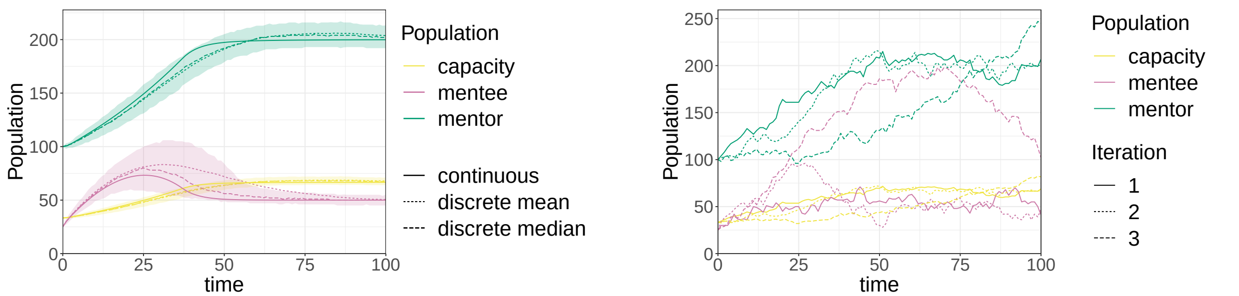

(a) Continuous and discrete models. (b) First three iterations of the discrete model.

Figure 1: Population dynamics.

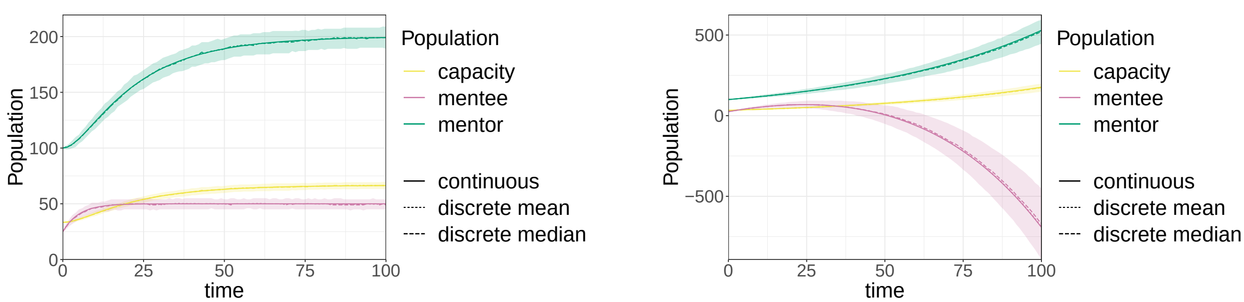

(a) $\min(x, ry)$ replaced with x. (b) $\min(x, ry)$ replaced with ry.

Figure 2: Comparison of continuous and discrete models with the non-linear term removed.

takes much longer to reach steady-state. This is noteworthy because both models are defined by the same dynamics. In particular, Figure 1b shows that significant excursions are possible like the third iteration example.

4 INTERPLAY BETWEEN NON-LINEAR AND DISCRETE EFFECTS

To further pinpoint the effects at play in Section 3, the benchmark scenario is simulated for both models without the non-linear term in Section 4.1, and at different population sizes in Section 4.2.

4.1 Non-linear Term

One hypothesis to explain what is causing the differences is that the non-linear term generates significant discrete effects at lower population sizes. To confirm this, the non-linear term is removed from the benchmark scenario in both models and the outputs generated by the models are then compared again. Figure 2 shows the effect of overlaying the continuous and discrete population curves after replacing $\min(x, ry)$ with only x or ry. This results in identical plots for the curves of the continuous model and the discrete model mean and median curves. Additionally, there is less spread in the discrete model and the spread is distributed evenly around the mean and median lines unlike in Figure 1a where the spread is biased upwards. This confirms that the non-linear term is the main contributor to the difference between the continuous and discrete models. In other words, given that only a single training regime is possible, the mean value of the discrete model is accurately predicted by the continuous model.

4.2 Low-Population Sizes

The population sizes of the benchmark scenario need to be verified if they are low enough to cause discrete effects, resulting in the differences between the continuous and discrete implementations. Consider the set of all models where the ratio between parameters is preserved with the scale adjusted such that $y_0 = 4x_0$ and $(x_{target}, y_{target}) = 2(x_0, y_0)$.

For the remainder of the paper, Δx refers to the change in mentee population for each event; in the discrete benchmark scenario this is 1. The change in mentor population will not be reported; this is done without loss of generality since a change in the mentor population for each event will always be equal to Δx. This is because both values are linked by the event where the mentee upgrades to a mentor. $\Delta x/x_0$ will be used to refer to the ratio between the change in population over the initial mentee population. x_0 is used as an analog for the total size of the population; the decision to use x_0 instead of y_0, x_{target}, or y_{target} is arbitrary since these variables are proportional to each other for the set of models under consideration.

The influence of discrete effects is examined by reducing the ratio of the change in mentee population at any given time step to the total population (i.e., making $\Delta x/x_0$ smaller). One way to think about this is that the discrete model is approaching the continuous model in the limit by reducing the time between events and the change in population at each event. This is done in two ways: scaling up the total initial and target populations, and recalculating the value for the intake (a) accordingly, or scaling down the change in population at each event and scaling up the chance of an event happening by the same factor. Note that both of these methods increase the number of events per time unit in the model to make up for the change in population at each event being a smaller proportion of the total population.

Figure 3 shows how the population dynamics of the discrete model are influenced when $\Delta x/x_0$ is reduced by a factor of 100. Figure 3a shows the mean, median and 25th-75th quantiles of the discrete model when initial and target populations are increased by a factor of 100 compared to the continuous model. Figure 3b displays the first three iterations of the discrete model with populations scaled by 100. Figure 3c and Figure 3d show the same results as Figure 3a and Figure 3b respectively, but instead of scaling up the population sizes by 100, the original benchmark populations are retained, and Δx is scaled down by 100. In Figure 3a, the effects of scaling up the initial and target population of the continuous and discrete models by 100 times show that the continuous model and the mean and median of the discrete model are now identical, and the ribbon showing the spread between the 25th and 75th quantiles is significantly smaller, implying less variation. The same result is seen in Figure 3c, where the populations are kept the same as in the benchmark scenario; however, the discrete change in population is reduced by a factor of 100. This implies that decreasing $\Delta x/x_0$ reduces the discrete effects of the model. This happens since decreasing $\Delta x/x_0$ by a factor of 100 results in 100 times more events with each event having 0.01 times the influence, which results in less variation. Figures 3b and 3d confirm this notion by showing the first three iterations of the simulations which are smooth and close to the mean, in contrast to the original discrete simulation of Figure 1b.

5 FACTORS AFFECTING THE DIFFERENCE BETWEEN MODELS

As shown in Section 4, the difference between the continuous and discrete models is caused by discrete effects from the non-linear term. However, the extent to which each factor affects the discrepancy is not known. Given the initial parameters x_0, y_0, x_{target}, y_{target}, a, b, c, and r, can it be determined whether the continuous model is viable without running any simulations? Furthermore, if the continuous model is not viable, can the difference between the continuous and discrete models be determined?

The system reaches a healthy steady-state when the training capacity is not exceeded, i.e., the system is in the unsaturated regime. The system variables are then bound by (3) which ensures that the rates for intake, absorption and attrition remain equal, while not exceeding the training capacity of the mentors. This means that changing any one of a, b or c by a given factor will result in changing all of these variables by that factor. Since the simulations treat the attrition c as given, and calculate a and b based on the values

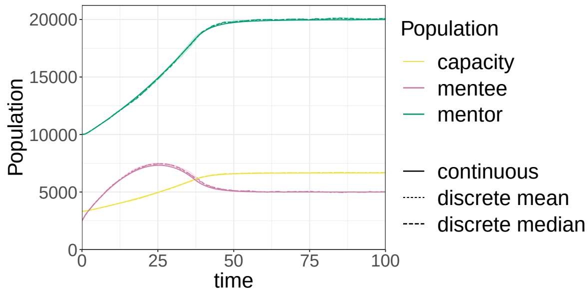

(a) Continuous and discrete models with population size scaled up by 100.

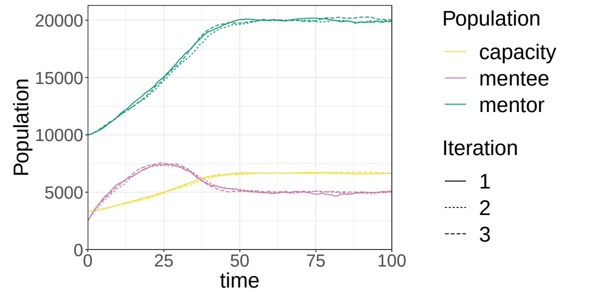

(b) First three iterations of the discrete model with populations scaled up by 100.

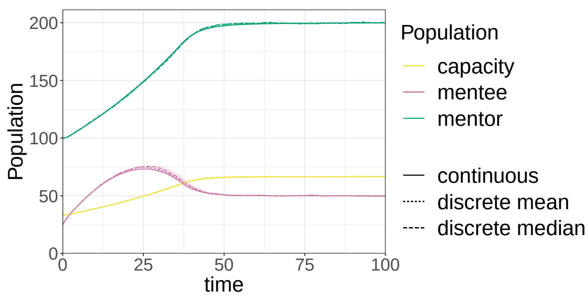

(c) Continuous and discrete models with Δx scaled down by 100.

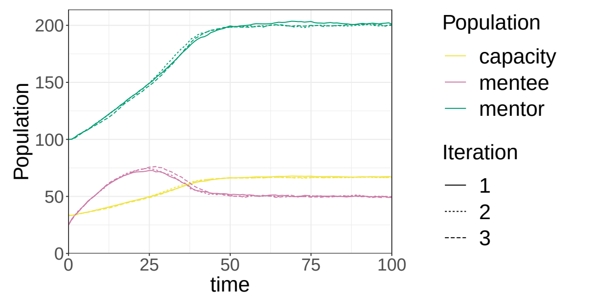

(d) First three iterations of the discrete model with Δx scaled down by 100.

Figure 3: Effects of reducing the magnitude of $\Delta x/x_0$.

of c, x_{target}, and y_{target}, the effect of changing c, while keeping the target populations constant is analysed. However, this has the effect of increasing/decreasing all the rates by the same factor which results in the time scale being stretched/compressed by that factor. This is not surprising; thus, we focus on the variables x_0, y_0, x_{target}, y_{target} and r.

Specifically, in Section 5.1, the effect of varying $\Delta x/x_0$ is considered in further depth, while, in Section 5.2, the concept of the saturation ratio is introduced and explored with an emphasis on how varying this ratio can affect the difference between continuous and discrete models is investigated.

5.1 Varying $\Delta x/x_0$

In Section 4.2, scaling $\Delta x/x_0$ down by a factor of 100 was shown to mitigate the discrete effects of the model. The next step is to determine what type of relationship exists between $\Delta x/x_0$, and the difference between the continuous and discrete models. To do this, $\Delta x/x_0$ must be varied and a measure must be chosen to effectively compare the continuous and discrete models.

There are two ways to vary $\Delta x/x_0$: iterate over values of Δx or over values of x_0. Both approaches should have the same effect such that the number of events and the influence of each event on the system will be the same *on average*. Therefore, Δx is chosen because the simulations will be simpler to compare if the initial and target populations are equal across all runs. Furthermore, the continuous model only has to be solved once (because it is deterministic), and all discrete model simulations will be compared to it. Using this approach, by reducing Δx and the time between events in the discrete model runs, the results should approach, in the limit, the output of the continuous model.

The chosen metric for comparing the continuous and discrete models is the average proportional difference (APD) between the continuous model's mentee curve (x_c) and the mean of the discrete model's

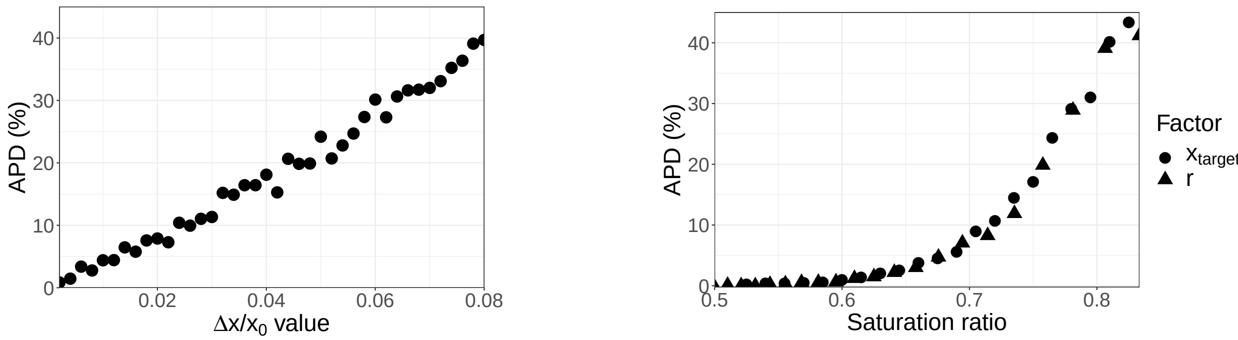

(a) Varying Δx.

(b) Varying saturation ratio.

Figure 4: APD between the continuous and discrete model mentee curves.

mentee curve ($\overline{x_d}$) and is given by (4), where n is the number of constant-time intervals at which each curve is sampled.

$$APD = \frac{n}{n-1} \sum_{t=0}^{n-1} \frac{|\overline{x_d} - x_c|}{x_c} \tag{4}$$

The mentee curves are also chosen because the difference between the two simulations is most noticeable when focusing on the mentee populations. The APD gives a metric which is invariable to scale, making is easier to compare two different simulations. This is especially important in Section 5.2 when the population sizes are different for each iteration.

The number of iterations for each data point increases as Δx increases; this is because each individual iteration takes significantly longer for smaller values of Δx given that significantly more events are generated. This choice is justified because the simulation has less variation as Δx decreases as shown in Figure 3 of Section 3. To obtain an accurate mean value for the Monte Carlo simulation when $\Delta x = 1.5$ may require thousands of iterations due to the large variation between runs; however, running thousands of iterations when $\Delta x = 0.01$ takes time on the order of weeks on a regular computer. The number of iterations (n_{iter}) for each Monte Carlo simulation is determined as follows:

$$n_{iter} = 1000\Delta x.$$

Figure 4a shows the results of taking the APD between the continuous and discrete models as Δx varies. It is clear that the difference between them is reduced as Δx is reduced. Figure 4a also quantifies the deviation due to the impact discrete effects on the dynamics when choosing the continuous model, a model computationally less expensive than the discrete model.

5.2 Varying the Saturation Ratio

Scaling x_{target} and y_{target} by the same factor will change the overall scale of the population. However, what change can be expected if x_{target} and y_{target} are scaled by different factors? In other words, what are the effects on the continuous and discrete models if the target ratio between mentees and mentors is not $1/4$, as is the case of the benchmark scenario? This target ratio is closely related to the capacity value r. If the target ratio is increased, then the steady-state mentee population will be closer to the mentor capacity. This leads to a new measure - the saturation ratio (s_r), which is calculated by dividing the target ratio by r. This is a metric that measures how close the steady-state populations are to the training capacity.

$$s_r = \frac{x_{target}}{r y_{target}}$$

To vary the saturation ratio, either the mentee or mentor target populations can be changed, or the value of r can be changed. Two simulations are run varying the saturation ratio from 0.5 to 0.85. The first simulation iterates through different values of r, and keeps all other variables the same. The second simulation iterates through target mentee populations, and b is recalculated for each iteration to achieve the steady-state population. At each iteration, both the continuous and discrete models are run, and then the APD between the mentee populations is calculated.

The results of the two simulations are shown in Figure 4b. The effects of varying the saturation ratio are similar whether you change r or the mentee target population. Furthermore, the trend shows that the lower saturation ratio values have less discrepancy between continuous and discrete models. The original model has a saturation ratio of 0.75, which results in the discrete model having roughly 17% more mentees, as seen in Figure 4b. It is also observed that the saturation ratio curve appears to taper off at higher saturation ratio values; this is likely due to the fact that the simulation does not reach steady-state in 100 time units, and thus does not have enough time to accumulate the total difference. Another important note is that when the saturation value dipped below ≈ 0.66, the continuous model remained in the unsaturated state for the entire run, while the discrete model had some iterations where the mentee population crossed into the saturated regime. This further shows how the continuous model and the discrete model diverge.

6 CONCLUSION

The aim of this paper was to devise a model of the mentee-mentor population dynamics where the absorption of mentees can saturate when there are not enough mentors. To explore these dynamics, both a deterministic continuous model and a stochastic discrete model were implemented. As demonstrated, even if a population starts and ends in a steady state with a healthy mentee-mentor ratio, a model that considers the system behaviour when the training capacity of the mentors is saturated should be used. Indeed, during the transition between the two states, the system may still enter several periods where the mentee-mentor ratio is high enough to restrict the absorption of mentees. Limitations of the continuous model when considering dynamics that switch between the unsaturated and saturated training regimes and at low-population sizes were also discussed. The degree to which these limitations affect the difference between continuous and discrete model implementations was further analyzed leading to the determination of a relationship between the total population size of the system and the difference between the continuous and discrete approaches. In particular, it was also observed that when the target mentee-mentor ratio approaches the latent absorption capacity, the difference between the continuous and discrete implementations was amplified.

For the benchmark scenario where the initial mentee and mentor populations are doubled, using a continuous approach to model the non-linear mentee-mentor dynamics instead of a discrete model was found to led to a significant under-estimation of the number of mentees in training and, consequently, the time to reach the new steady state. Thus, using a continuous model of the mentee-mentor population dynamics to plan a transition scenario should be made circumspectly and fully aware of its limitations, especially in terms of the estimation of the time to complete the transition, which is usually of significance for military planners.

REFERENCES

Bastian, N. D., and A. O. Hall. 2020. "Military Workforce Planning and Manpower Modeling". In *Handbook of Military and Defense Operations Research* (1st ed.)., edited by N. M. Scala and J. P. Howard. CRC Press.

Boileau, M. L. A. 2012. "Workforce Modelling Tools Used by the Canadian Forces". 18–23.

Bryce, R., and J. Henderson. 2020, 12. "Workforce Populations: Empirical versus Markovian Dynamics". In *2020 Winter Simulation Conference (WSC)*, 1983–1993.

Diener, R. 2018, 01. "A solvable model of hierarchical workforces employed by the Canadian armed forces". *Military Operations Research* 23:47–57.

Forrester, J. 1965. *Industrial Dynamics*. System dynamics series. Pegasus Communications.

Henderson, J., and R. Bryce. 2019, 12. "Verification Methodology for Discrete Event Simulation Models of Personnel in the Canadian Armed Forces". In *2019 Winter Simulation Conference (WSC)*, 2479–2490.

Masuda, N., and L. E. C. Rocha. 2020. *A Gillespie algorithm for non-Markovian stochastic processes*.

Mehta, P. 2019. "Random walks, Birth Death Processes, and the Gillespie Algorithm.".

Novak, A., L. Tracey, V. Nguyen, M. Johnstone, V. Le, and D. Creighton. 2015. "Evaluation of tender solutions for Aviation Training using discrete event simulation and best performance criteria". In *2015 Winter Simulation Conference (WSC)*, 2680–2691.

Okazawa, S. 2020, 12. "Methods for Estimating Incidence Rates and Predicting Incident Numbers in Military Populations". In *2020 Winter Simulation Conference (WSC)*, 1994–2005.

Séguin, R. 2015. "Parsim, a simulation model of the royal canadian air force (rcaf) pilot occupation - an assessment of the pilot occupation sustainability under high student production and reduced flying rates". In *Proceedings of the 4th International Conference on Operatinal Reserach and Enterprise Systems (ICORES)*.

Suvorova, S., A. Novak, B. Moran, and T. Caelli. 2019. "The Use of Markov Decision Processes for Australian Naval Aviation Training Schedules". *Military Operations Research* 24(2):31–46.

Swift, R. J. 2002. "A Stochastic Predator-Prey Model". In *Bulletin of the Irish Mathematical Society*, Number 48. Irish Mathematical Society.

Vincent, E., and S. Okazawa. 2019. "Determining Equilibrium Staffing Flows in the Canadian Department of National Defence Public Servant Workforce". In *2019 International Conference on Operations Research and Enterprise Systems (ICORES)*, 205–212.

Zais, M., and D. Zhang. 2015, 11. "A Markov chain model of military personnel dynamics". *International Journal of Production Research* 54:1–23.

AUTHOR BIOGRAPHIES

SAMUEL SCHAFFEL is a research intern at Defence Research and Development Canada – Centre for Operational Research and Analysis. He is working on completing his Bachelor of Applied Science and Engineering in Applied Mathematics with a focus on systems and robotics at Queen's University in Kingston, ON, Canada. His e-mail address is samuel.schaffel@forces.gc.ca.

FRANÇOIS-ALEX BOURQUE is a Defence Scientist at the Defence Research Development Canada – Centre for Operational Research and Analysis. He was a Senior Scientist at the NATO Centre for Maritime Research and Experimentation from 2015-2018 and a Visiting Scientist in 2013. Dr. Bourque holds a PhD in theoretical nuclear physics from McGill University. As a Defence Scientist, he has provided decision support to the Canadian Department of National Defence on a variety of projects. Additionally, he has explored the potential defence and security use of unmanned systems. Dr. Bourque can be contacted via email: alex.bourque@forces.gc.ca.

SLAWOMIR WESOLKOWSKI is a Senior Engineer in Department of National Defence who provides advice on analytics to the Government of Canada. Dr. Wesolkowski worked previously at Defence Research Development Canada – Centre for Operational Research and Analysis, NCR Canada Ltd., and various other companies. He holds PhD, MASc and BASc degrees in Systems Design Engineering from the University of Waterloo. He is an INFORMS Certified Analytics Professional (CAP®) and the recipient of numerous awards for analytics including the 2018 Royal Canadian Air Force Commander Commendation. His email address is slawomir.wesolkowski@forces.gc.ca.