“Learnoids”: One Million Synthetic Students

Competencies Model Generation, Simulation and Results

FINAL DRAFT FOR LIMITED CIRCULATION

Marc Levesque, Cuiting Li, Carnegie-Mellon University (CMU), METALS program
Charles Fadel, Robbie Taylor, Center for Curriculum Redesign (CCR)

August 1, 2022

Made possible by a generous grant from the Koshland Foundation

Abstract

The Center for Curriculum Redesign’s Framework details 12 Competencies and 60 Subcompetencies which collectively describe the most important skills, character qualities, and meta-learning abilities for people in the 21st century. The work described in this paper is the fruit of a collaboration between CCR and CMU, through a generous grant from the Koshland Foundation. It establishes methods for ways that progression of the Competencies may be simulated and how different personality types and influences such as school environments and parenting impact student progress. While these variables were utilized to hypothesize the influencers of the model, any variables could be integrated into the model. Our model establishes the potential for being able to detect not only if the student is experiencing negative situations, but that with enough data it may be possible to detect exactly what the student is experiencing, giving teachers, parents, and others the ability to assist the student with relative precision. It can also be determined how students with similar personalities progress, which students may be the most at risk or the most naturally capable in which Competencies, and what sorts of interventions and support can best work with students both as individuals and in groups. Lastly, it may be possible to infer theories of mind based on aggregate student data, showing for instance what developmental stages accompany various competencies, their interrelatedness, etc.

This approach to generate synthetic students mirrors Data Science techniques in Health Care, where a pilot size may be small (<50 patients) yet the models need to analyze the broader impact, and hence generate (>50,000) “synthetic patients”. And as in Health Care, this approach offers the future, tantalizing possibility of “Digital Twins” to simulate and test interventions before they are deployed to the student. With this proof of concept, real student data would now build opportunities to further refine the model and create a framework for reporting on and building student competencies.

CENTER FOR CURRICULUM REDESIGN
Making Education More Relevant

Part 1: Introduction

The student progression simulation work aims to create one million synthetic students to showcase the power of Learning Engineering to provide student-level information for interventions, and group-level information leading to theories of mind. It is based on profiles of a number of real students in order to accomplish the three following objectives:

1. Be able to determine how to find the most positive interventions that teachers can do, such as, for instance, by focusing on Metacognition-targeted interventions and observing its correlated effects on the other competencies.
2. Detect the malleability of students, such as determining how much a positive school environment can truly impact a student's competency scores.
3. Be able to detect through an algorithm whether a student is undergoing a negative effect, such as bullying, at a given point in time.

We do remind the reader to keep in mind that this is merely a simulation, and to not mistake the actual results for reality.

Establishing context

We began our research by studying which of the competencies are most commonly found to be correlated. We also researched how the individual competencies tend to evolve over time and what factors are commonly known to influence them, either in beneficial or detrimental ways.

Trends, Correlations & Anti-Correlations Research

In order to determine how the competencies might be correlated and how they could be expected to progress over time, we researched each possible pairing of the competencies (see Trends, Correlations & Anti Correlations in Appendix A). For example, we researched Creativity paired with each of the 11 competencies, and then Critical Thinking with the 10 competencies other than Creativity, and so on. However, there are many competencies that are frequently observed to be well-developed when certain other competencies are well-developed, implying correlation if not causation. For example, Mindfulness and Metacognition are typically both at similar levels, and it would be unusual to see a case where Metacognition is very high but Mindfulness is very low. There are also many cases where there seems to be a slight positive link, and some where we could not find any link between the pairs. The only instance where we found a negative relationship is where extreme Courage, specifically expressed as overconfidence, is shown to be linked to low Growth Mindset.

We also researched how the competencies would likely progress on their own. Generally, they tend to increase slightly over time, with the exceptions of Curiosity and

Growth Mindset which may decrease during adulthood, and Creativity which seems to increase until around 4th grade, decrease until 6th, and then increase for at least a few more years.
All of these perceived correlations and trends then became the basis for our model, specifically as we modeled that targeted increases in certain competencies will have a lesser or greater effect on certain other competencies.

Outside Influences

We also researched which outside factors might have a detrimental effect on student progress on the competencies (see Key Outside Influences in Appendix A). Since part of the purpose of measuring the competencies in the first place is to know if a student is not making as much progress as they should, it is important to be aware of some reasons for why students might have slower-than-normal or even negative changes over time. Needless to say, there are numerous reasons for why a student might regress, but outside influences that we can especially tie to competencies include bullying, which is mediated by Resilience, low socio-economic-status which leads to slower progress on the competencies in general, and so on.

Malleability

Further research was done on the topic of malleability, particularly aiming to discover which competencies teachers will most easily be able to influence (see Malleability in Appendix A). While Growth Mindset seems a promising candidate, there were no competencies that we found to be exceptionally malleable, *at least via the very limited research done for this simulation*. A likely bigger factor on malleability is the degree to which a student's personality influences their competency progression, which will be discussed later in this paper.

Personalities

We also considered how students' personalities might naturally influence their progression on competencies, regardless of outside influences. We chose to focus primarily on the Big 5 personality types because of its global availability. We initially tried creating student personas based on these personality types, but instead found the best method was to ask teachers to create student personas based on their own students and to then match those personas with the characteristics of the Big 5.

Data Collection

In order to determine baseline levels for student scores, on the 1-10 lifetime scale, we surveyed 5 teachers and interviewed 2 others (see Personas in Appendix A). To do so, we informed teachers that they were to:

- Imagine students aged 10
- Come up with 3 contrasting students
- Assign the students' scores of 1-10 on each of the 12 competencies, keeping in mind that a student of that age would likely score in the 3-5 range

Once the teachers assigned the scores, we subtracted 2 from each score so that we could create a potential baseline for a 6-year-old, as we intended to show the progression from age 6-18, which are the typical 1st-grade to high-school years.

Part 2: Model Creation

The spreadsheet with the baseline formulas, factors, and 21 personas and a number of simulated score scenarios is available is titled Model Progression (see Appendix A) but would be best viewed after reading through the notes below:

Base Formula and Scale

We knew we wanted the possible score to range from 1 to 10 over a student's lifetime, and that it should be more difficult to make progress over time. Therefore, we also added in a measure of "opportunity points," which could be described as naturally arising or deliberately assigned opportunities for the students to make progress on the competencies. The difficulty curve is inspired by the points progression system commonly found in video games progression between skill levels that requires increasingly large numbers of opportunity or experience points.

The opportunity point scale ranges from 1 to 100, and, using a diminishing returns curve (diagram below), we placed the midpoint, 50, at 7 on the 1-10 scale. This means it is as easy (or difficult) to go from 1-7 on the 10 point scale as it is from 8-10.

We were then able to plot out the relationship between the opportunity points and the score and derive a formula which we can use to convert to and from opportunity points and scores.

We created the following terminology to assist with converting between the 10-point and 100-point scales:

Starting Score (from 1-10) = SS
Start Points (from 1-100) = SP
Final Points (from 1-100) = FP
Final Score (from 1-10) = FS

The formula is then as follows:

Converting SS into SP: $SP = 0.95*SS^2 + 0.5*SS$
Converting FP into FS: $FS = \sqrt{76*FP+5}*0.117688-0.263158$

To interpret on a scale of 1-4, we can simply divide the 10 point scale evenly by 4 so that 0-2.5 points is 1, 2.5-5 points is 2, and so on. Either way, the number of points needed to reach a higher score increases over time, as seen by this chart:

In order to imagine how a student might naturally progress over time, we allotted them an average of 3 opportunity points worth of progression for each competency per year, starting at a score of 2 and opportunity points of 5 for a 6 year old. Thus, a 7 year old who had one year of practice will on average go up 3 opportunity points, to 8, which translates into a score of 2.65. By the time they're 18, they'll have had 41 opportunity points and a score around 6.3.

To explore how progression might continue after age 18, for most competencies we modeled that the brain malleability of the student might lessen over time, so they are less likely to make personality/competency changes easily. Thus, past 18, instead of progressing by 3 opportunity points each year, they'll progress by 2.5, then 2, then 1.5, and so on. Following this trajectory, having 41 opportunity points and a score of 6.3 is realistic around age 18, but reaching a perfect 100 points and score of 10 in **all** competencies might take upwards of 190 years!

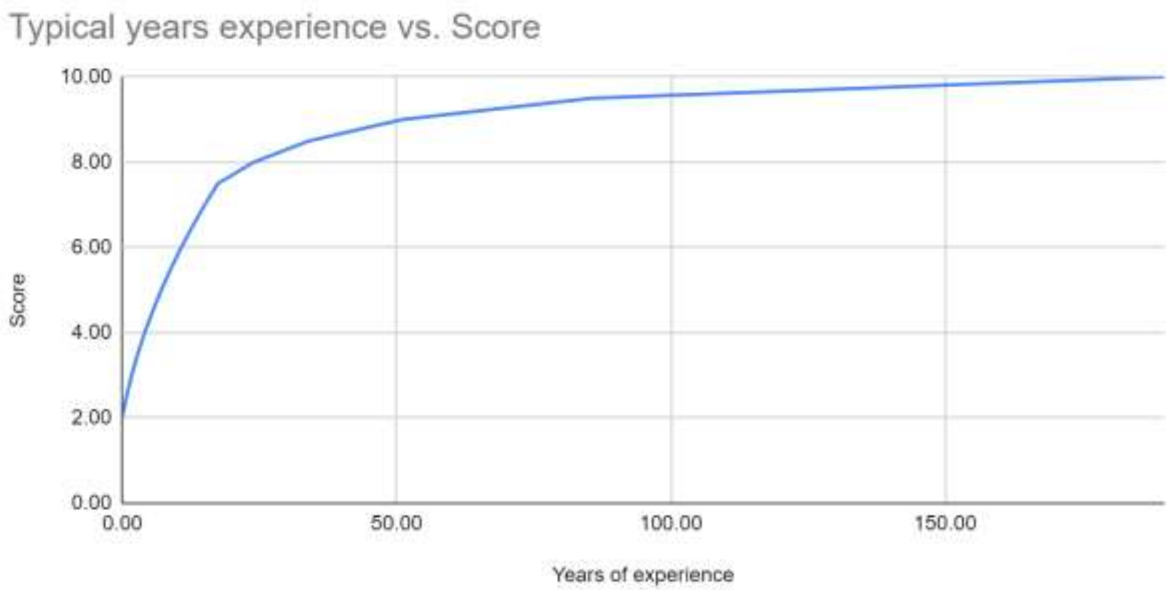

It is the axiom of this paper that deliberate practice yields development of competencies. As such, some competencies such as leadership likely have numerous practice opportunities in professional settings, and so may have a lower rate of slowing, if not faster progress, after age 18.

This form of progression is sufficient for now, but for future research, it would be worth determining typical progression speeds and curves for each competency, some of which may have an s-shaped, n-shaped or other type of curve rather than our current model.

Factors

We then included several possible factors that might alter how students progress. The idea is that the student will start out with a score from 1-10, then that will be converted to the 100-point scale, then they will experience some degree of change over time from the various factors, and then at whichever end point we are observing, in this case when the students complete high school at 18, the points are converted back into the 1-10 scale. Various combinations of factors are added so that each possible outcome is modeled. The conversion can happen at any point though along the way if more granularity is desired.

The factors that we currently are modeling are school environment, personality, parental involvement, bullying, and frequent school switching. Those will all be described shortly in more detail, but similar modeling could be done for any factor that might affect students.

School Environment

We imagined that some students likely go to schools that constantly encourage and support them and give them lots of opportunities to practice their competencies, and some go to schools that offer medium or poor levels of support. For the purposes of the model, we assumed that all children will attend school, although "school" could be reinterpreted to mean homeschooling or even building life experience without formal study, in which case the parameters and curves will need to be adjusted accordingly.

For the most baseline level in a high-performing school, in order to determine how a student would end up on each competency from the age of 6 to 18, we decided that each competency would be a primary focus for one third of each year, so each competency is a primary focus once every 4 years. The primary focus competencies increase by 4 points, two closely related competencies increase by 1, most of the rest increase by 0.7, and two hardly related competencies increase by 0.55. This works out so that the average increase across all competencies from age 6-18 is in line with what we would expect.

Based on our research on progression on the individual competencies over time, we placed a negative weight on Creativity from ages 10-12 as a drop is commonly observed there, and we slightly reduced the overall progress of Growth Mindset as that is also shown to not typically increase much without any intervention over time.

To create a more neutral school representation, we divided the high-performing school increases by 2, so a 4 became a 2, a 0.7 became a 0.35, and so on.

To create a representation of a low-performing school, we divided the high-performing school increases by 3, so a 4 became a 1.33, and a 0.7 became a 0.33, and so on.

In most cases, this means that the student in the low-performing school environment will make significantly less progress from age 6-18 as compared to a student in a low-performing school.

Personality

We grouped each of the 21 personas created by the teachers into one of the five Big 5 personality types and added some positive and negative weight to some competencies that we felt aligned with the descriptions by the teachers and by the personality types.

Extraversion: We determined that communication, collaboration, and leadership are positively associated with Extraversion, and Mindfulness may be negatively correlated. Therefore, we assigned a natural 0.2 opportunity point increase per 4 months for Communication, Collaboration, and Leadership, and a -0.2 decrease for Mindfulness.

Openness: We placed 0.2 positive weights for Creativity, Critical Thinking, and Communication, and a negative weight for Mindfulness.

Conscientiousness: We placed 0.2 positive weights for Mindfulness, Metacognition, and Growth Mindset, and -0.2 weights for Creativity and Leadership.

Neuroticism: Students with this personality type seemed to struggle strongly over time, so we placed -0.2 weights on Mindfulness, Resilience, Metacognition, and Growth Mindset with no positive weights.

Agreeableness: We placed 0.2 weights on Collaboration and Ethics, and -0.2 weights on Creativity and Critical Thinking.

See Table 1 in Appendix B for the resulting scores.

Since we do not yet know exactly how much a given student's personality will affect their competency progress, we developed a combination where personality has a much larger influence on the student's overall development, greater than that of their schooling.

To create this result, we multiplied the original personality influences by 1.5 and added 10 to each. We then divided all the other factors, such as school and parental influence, by 2, resulting in students generally reaching a similar stage at the end of highschool as we predicted, but having gotten to that stage more so as a result of their personality or internal factors rather than external factors. See Table 2 in Appendix B for resulting scores.

Parental Involvement

We figured that great parental involvement might result in a general increase in most competencies of 0.2 per 4-month period, but a -0.2 change in Resilience as the students may be more sheltered. In addition, we modeled that Metacognition and Growth Mindset might have particularly strong positive weights, in this case 0.8, for the first 4 years, then 0.5 for 4 years, then 0.3 for 4 years as the student is less and less affected by how they interact with their parents.

For negative parental involvement, we placed a -0.4 weight on most competencies for the students' first 8 years, but a positive 0.2 impact on Resilience as the student is forced to build resilience, and we left Ethics unchanged as we did not find much of an impact on that based on parental involvement.

Negative Factors

Students are often bullied in Middle School, so we modeled what that might look like. They will perhaps experience a -0.4 change in all competencies for two years that they are being bullied, and then have lingering effects of -0.1 on all other competencies for 2 additional years.

Students who switch schools often might also have negative effects. For a student who, in addition to the regular transitions, switches schools twice during grades 1-5 and once more during High School, we placed a 16 month -0.2 per month effect on most sub competencies. We also placed a -3.0 per month for 4 months effect on Courage, Resilience, and Growth Mindset, a -0.7 effect each month for 4 additional months, -0.5 each for 4 more, and then -0.2 per month for the remainder of the 16-month span, thinking that those competencies are the ones that may be the most strongly affected by having to switch often between schools, lose friend groups, and so on.

Minor Correlations

Minor correlations naturally occur in the model due to our arrangement, discussed earlier, where we set large increases in certain competencies to have lesser or greater effects on the other competencies based on how closely related we consider them to be.

Major Correlations

Along the same lines, we also modeled what might happen if the competencies are even more dramatically correlated. To do this, we pulled from CCR's work on how often certain competencies are mentioned in papers about other competencies, seen here:

	Creativity	Critical Thinking	Communication	Collaboration	Mindfulness	Curiosity	Courage	Resilience	Ethics	Leadership	Metacognition	Growth Mindset
Creativity	72%	13%	0%	0%	0%	5%	0%	0%	0%	0%	10%	0%
Critical Thinking	14%	68%	0%	0%	0%	4%	0%	0%	0%	0%	14%	0%
Communication	0%	0%	77%	6%	0%	0%	1%	2%	2%	6%	7%	0%
Collaboration	0%	0%	21%	53%	0%	0%	0%	0%	1%	22%	3%	0%
Mindfulness	0%	0%	0%	0%	81%	0%	0%	9%	0%	0%	10%	0%
Curiosity	5%	3%	0%	0%	0%	77%	1%	1%	0%	0%	11%	2%
Courage	0%	0%	4%	1%	0%	1%	49%	39%	0%	1%	5%	1%
Resilience	0%	0%	5%	0%	3%	1%	20%	59%	1%	1%	10%	0%
Ethics	0%	0%	5%	1%	0%	0%	0%	1%	72%	7%	14%	0%
Leadership	0%	0%	12%	14%	0%	0%	0%	1%	6%	61%	6%	0%
Metacognition	2%	2%	6%	1%	1%	2%	1%	4%	5%	3%	74%	0%
Growth Mindset	0%	0%	0%	0%	0%	4%	1%	1%	0%	0%	1%	94%

For example, across several papers on Creativity, Critical Thinking was mentioned 14% of the time. This does not necessarily mean that there is a 14% correlation between Creativity and Critical Thinking, but it suggests that there is some sort of link and is a good starting point for creating more distinct correlations.

Noise

We experimented with several different levels of noise in order to determine how our model would perform with a given level of noise. Ultimately, we most frequently used a 20% noise level. This means that any improvement on any competency will vary by 20%, and more specifically that the 20% variation is 1 standard deviation.

Targeted Teacher Interventions

The final factor we considered are targeted positive interventions led by teachers. Specifically, we modeled the following interventions:

- A “Skill Dimension” intervention where teachers get 36 additional opportunity points and spend them on the Creativity, Critical Thinking, Communication, and Collaboration competencies, and the other competencies then vary based on how they are correlated to the Skill Dimension targeted competencies.
- A “Character Dimension” which worked the same way but targeting the Leadership, Ethics, Resilience, Courage, and Curiosity competencies.
- A “Meta-Learning Dimension” which worked the same way but targeted the Growth Mindset, Metacognition, and Mindfulness competencies.

Part 3: Simulation Process

Once we could calculate how much to add or subtract to a given student in the case of any combination of the above factors, we were able to then synthesize large sets of students along with their unique corresponding competency scores. See the Github link Appendix A for sample code.

The typical simulation process is as follows:

Within a spreadsheet:

1. A baseline will be assigned using the 21 teacher-generated personas that span across the Big 5 personalities.
2. Each of those personas will be combined with a combination of one or more outside factors, such as a 3-point decrease on various competencies due to bullying.
3. This process will be repeated until most or all possible combinations are covered.
4. The result is then rows of competency and persona combinations, usually about 120 in total, that include a label for the combination type, a starting score for

each competency, an ending score based on the selected combination of factors, and the noise level by which the scores should vary.

In R:

1. The rows of information then can be fed into an algorithm in R which will take in each of these values, add the noise, multiply by an assigned value to reach the desired sample size, and return either the 10- or 100-point scale results in a table that can be then exported as a spreadsheet in order to be used for machine learning.

Part 4: Simulation Results

Once we had a model that allowed us to generate synthetic students, we were able to use machine learning to determine what we would be able to discover in the data, which machine learning algorithms worked best, and with what degree of accuracy we could correctly make predictions about student progress.

Sufficient Differences for Detection

In order to examine how different the data would need to be in order for us to correctly categorize students with any degree of accuracy, we generated several data sets that contained increasingly distinct sets of data.

For example, one data set had 5,000 students who all progressed through K-12 the same way except that some of the students had a score on a single competency that was 20% lower than the rest. In another set, there were students who were 20% lower on all 12 competencies than the rest of the students. Overall, we found that if all 12 competencies vary from standard progress by about 0.3 to 0.5 times the amount of noise, which in this case we set at 20%, then we can correctly categorize students at least 70% of the time. If only one competency is varied, for example if some students are progressing normally and some are nearly the same but have a low Creativity score, that score will need to on average be different from the baseline by an amount equal to the noise in order to achieve 70% accuracy.

Similarly, if there are three different outcomes, such as a student progressing normally, a student progressing slower than usual, and one progressing faster than usual, the scores should vary on at least 6 out of 12 competencies by an amount equal to the noise level in order to be accurate in identifying the students.

If the scores are considered only over the course of a single year, rather than all the way through K-12, then the best results came from predicting purely if the student was undergoing a negative event or not, in this case with 83% accuracy. If more fidelity is desired, then it is necessary to make use of historical data for that student so that the model can better differentiate between outcomes.

Points of Confusion Among Scenarios

We also examined which outcomes were most confusing for the model. Naturally, the more similar the outcome the lower the accuracy, but in our model that particularly turned out to most often be negative events being confused with a lack of positive events. Take for example, the following confusion matrix. The labels are “GP” meaning good parenting, “BP” meaning bad parenting, “NS” meaning normal or average performing school, “LS” meaning low-performing school, “HS” meaning high-performing school, and “NE” meaning negative event, in this case either bullying, school switching, or a combination of the two.

a	b	c	d	e	f	g	h	i	j	k	l	<-- classified as
351	0	0	0	0	0	1144	0	2	503	0	500	a = .GP.NS
9	2240	0	0	1	0	0	228	0	0	22	0	b = .BP.NS
0	0	1429	4	0	0	0	0	1067	0	0	0	c = .GP.HS
1	0	0	1613	0	0	0	0	133	753	0	0	d = .BP.HS
12	4	0	0	1085	0	1145	0	0	0	254	0	e = .GP.LS
0	0	0	0	0	1727	0	702	0	0	0	71	f = .BP.LS
110	1	0	0	635	0	6112	0	0	55	587	0	g = .GP.NS.NE
2	587	0	0	0	792	0	5856	0	0	6	257	h = .BP.NS.NE
3	0	768	119	0	0	0	0	6593	17	0	0	i = .GP.HS.NE
26	2	0	720	0	0	36	0	13	6703	0	0	j = .BP.HS.NE
9	0	0	0	489	0	195	50	0	0	6757	0	k = .GP.LS.NE
0	0	0	0	0	186	0	8	0	0	0	7306	l = .BP.LS.NE

In the table, it can be seen that the biggest point of confusion is in column g row e, *students who experience good parenting but are in a low-performing school are being confused by the model with students who are experiencing good parenting in an average school but are undergoing a negative event.* Fortunately, this confusion can be addressed in the future by knowing more accurately which competencies are affected by the negative events, therefore allowing the model to separate between an average school with negative events and a low-performing school with no negative events. In the short term, *this confusion can also be reduced by having a human take a look at the possible outcomes for a given student and by taking other contexts into accoun*t. For example, if the school on average is known to be an average school, then the result with the average school and negative event outcome would be the more likely outcome rather than the low-performing school and no negative event. Similar distinctions could be made by teachers who know their students well. For example, if a teacher knows if a student more likely has a good or bad parental influence, for example, then the teacher can take that influence into consideration when considering if the student is having a good parenting with no positive intervention result, or a negative parenting with a positive intervention result, which may look similar to the algorithm but will be distinct to the teacher.

Predicting 24 Outcomes

Once we knew what degree of difference was typically necessary for detection, and which events were the most confusing, we expanded the original 21 personas to 120 types of students and generated over 1 million synthetic students based on those 120 types. We kept noise at 20%.

We found that if we had a good idea of what the outcomes might be, such as if we were then trying to recover the same 120 outcomes, then simple algorithms such as Naive Bayes were highly effective, typically at least 70%. If we hypothesized that we would have less clear of an idea of the possible types of outcomes and instead group the results into 24 larger categories, then using a more complex model such as Support Vector Machines or Multilayer Perceptrons were somewhat more effective, for example SMO reaching 70% accuracy whereas Naive Bayes reached 68%. As the number of scenarios we looked for decreased, the differences grew between Naive Bayes and Multilayer Perceptrons in particular, with the latter performing relatively better.

Personality and Prediction Accuracy

For the earlier predictions, we used a low-impact personality model, where personality accounted for roughly 25% of a student's natural progression, and the other events such as schooling and parenting accounted for the remainder. *However, we also experimented with a version where personality accounted for 50% or more of the student progress.*

We found that this version generally meant that outcomes were harder to predict. This is because, as the personality becomes a larger contributing factor, the differences caused by schooling, parenting, negative events, and so on get increasingly small and therefore difficult to tell apart.

Assuming we don't know how the student's personalities are affecting the 50% of their progression, then we found that we could predict 12 major outcomes with about 65% accuracy using a Multilayer Perceptron with 1 million students. If we aimed only to predict whether the student was progressing normally or experiencing a negative event, then Naive Bayes gave us 75% accuracy and Multilayer Perceptron gave us 85%.

This suggests that, depending on the degree to which a given student's personality affects their progress, it will be helpful to have a clear understanding of how their personality impacts their scores so that that impact can then be separated from the task of predicting and categorizing based on the other progression factors. With real data and a better understanding of the relative influences of competency development, the model can be refined.

Personality and Necessary Intervention

For each of the Big 5 personalities, our model simulated different levels of progression on each of the competencies. We found that, assuming that personality made up 25% of student progress, then each student and personality category would likely progress similarly to other students and personality types since they would ultimately end up being exposed to similar influences such as similar school environments which then acted as an equalizer (see Table 3 in Appendix B). However, if personality was 50% or if school for any reason did not have much effect, then the personalities would diverge in how they progressed. For example, in the following charts, we have listed a starting and ending score for Metacognition from age 6 to 18 for each of the Big 5 personalities. Even if the personality is the strongest influence, the Conscientiousness student seems to make similar progress to a situation where the school is the stronger influence. However, the Openness and Extraversion personalities both perform worse on Metacognition if their personality is the strongest factor, and the Neuroticism personality is at risk of a drastically lower score if the school influence is not strong.

While this is all theoretical, it raises many interesting questions about how *school interventions may need to take into account student personality and other factors in order to provide the right level of support for each student.* As one-size-fits-all approaches are often ineffective, these types of data sets can equip decision-makers with useful information to guide support programs and interventions for learners.

Correlations and Intervention Impact

Another important area of study is which interventions are the most effective, even if given limited resources.

As mentioned earlier, we created two different levels of correlation between the competencies. In one level, each competency is only mildly correlated and therefore only mildly affects the other competencies. In the other level, the competencies are strongly correlated and Metacognition is particularly likely to affect the progression of the competencies.

To determine how these correlations might affect intervention effectiveness, we modeled an intervention where a teacher has 36 points that they can put towards helping their students. In our early modeling work, we had modeled that bullying would cause a 3 point decrease per year per competency for students, so 36 points is enough to undo the effects of one year of bullying if the points are spread out to 3 points each for all 12 competencies.

Spreading out the points among all 12 competencies would indeed seem a reasonable thing to do in order to counteract a generally negative event or give an extra boost to students, but we imagined that the teacher also has one other option: They could try a Meta-dimension intervention where, instead of putting 3 points on each competency, they put 12 points each on Metacognition, Growth Mindset, and Mindfulness.

Our simulation result, created with over 1 million students for each option, reveals that if there is no strong correlation between the competencies, then it hardly matters how the points/intervention is distributed. However, if there is a correlation between the competencies, and if Metacognition is particularly influential as we suspect, then focusing just on the three competencies in the Meta-dimension results in a substantial improvement with an average. In our simulation (see Table 4 in Appendix B), we observe that the presence of a strong correlation with a Meta-dimension intervention results in an overall average score of 44 points vs evenly distributing the points and only achieving a 39 point average. This roughly 10% difference is enormous and would be achieved by knowing how the competencies are correlated and therefore knowing which competencies, when boosted, have great impact on the rest.

Part 5: Conclusions

While the data is for now synthetic, these efforts demonstrate a small part of the enormous potential of applying similar analysis techniques to real student data. There is the potential for being able to, while measuring student progress in the Competencies, detect not only if the student is experiencing negative situations, but with enough data it may be possible to detect exactly what the student is experiencing, giving teachers, parents, and others the ability to assist the student with personalized interventions. It can also eventually be determined how much a student's personality affects their progress, how students with similar personalities progress, and what sorts of interventions and support can best work with those individual students. This offers the tantalizing future possibility of "Digital Twins" to test interventions before affecting the real student.

With more analysis of the correlations between the competencies, teachers with limited time and resources can get recommendations as to the most impactful interventions that they can do with their class, and those well-targeted interventions can have a compounded effect on the correlated competencies and boost student performance far beyond what would be possible with a general less targeted less data-informed intervention.

As more measurements are made, and as more data is collected and analyzed and modeled, accuracy and efficacy of the predictions, identification, and interventions will only increase.

Appendix A: Data Sources Links

(If inactive or unreachable, please email CCR at info@curriculumredesign.org)

Key Outside Influences
Personas (teacher generated)
Trends, Correlations & Anti Correlations
Simulation Code - Github Link
UX demo website

Appendix B: Tables

Table 1:

Personality	Creativity	Critical Thinking	Communication	Collaboration	Mindfulness	Curiosity	Courage	Resilience	Ethics	Leadership	Metacognition	Growth Mindset
None	0	0	0	0	0	0	0	0	0	0	0	0
Extraversion	0	0	7.4	7.4	-7.4	0	0	0	0	7.4	0	0
Openness	7.4	7.4	7.4	0	-7.4	0	0	0	0	0	0	0
Conscientiousness	-7.4	0	0	0	7.4	0	0	0	0	-7.4	7.4	7.4
Neuroticism	0	0	0	0	-7.4	0	0	-7.4	0	0	-7.4	-7.4
Agreeableness	-7.4	-7.4	0	7.4	0	0	0	0	7.4	0	0	0

Table 2:

Large Personality	Creativity	Critical Thinking	Communication	Collaboration	Mindfulness	Curiosity	Courage	Resilience	Ethics	Leadership	Metacognition	Growth Mindset
None	10	10	10	10	10	10	10	10	10	10	10	10
Extraversion	10	10	21.1	21.1	-1.1	10	10	10	10	21.1	10	10
Openness	21.1	21.1	21.1	10	-1.1	10	10	10	10	10	10	10
Conscientiousness	-1.1	10	10	10	21.1	10	10	10	10	-1.1	21.1	21.1
Neuroticism	10	10	10	10	-1.1	10	10	-1.1	10	10	-1.1	-1.1
Agreeableness	-1.1	-1.1	10	21.1	10	10	10	10	21.1	10	10	10

Table 3: Competency Performance with Strong Personality vs Strong School Influences

Openness

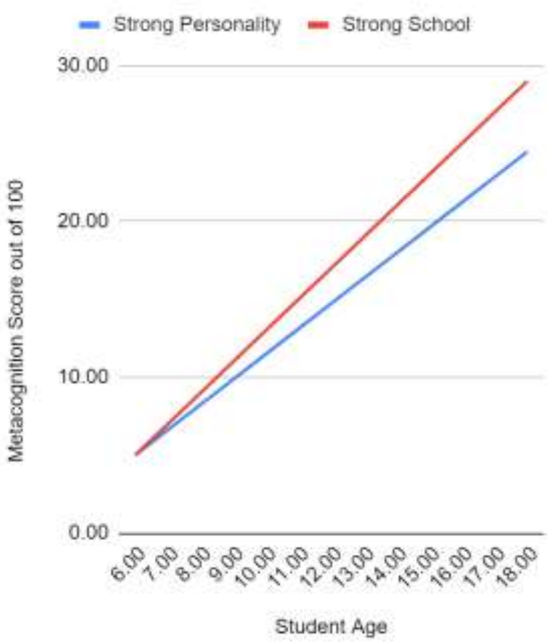

Extraversion

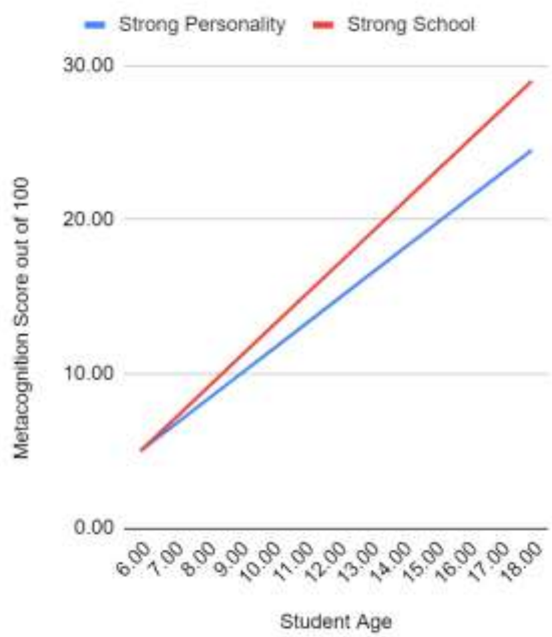

Conscientiousness

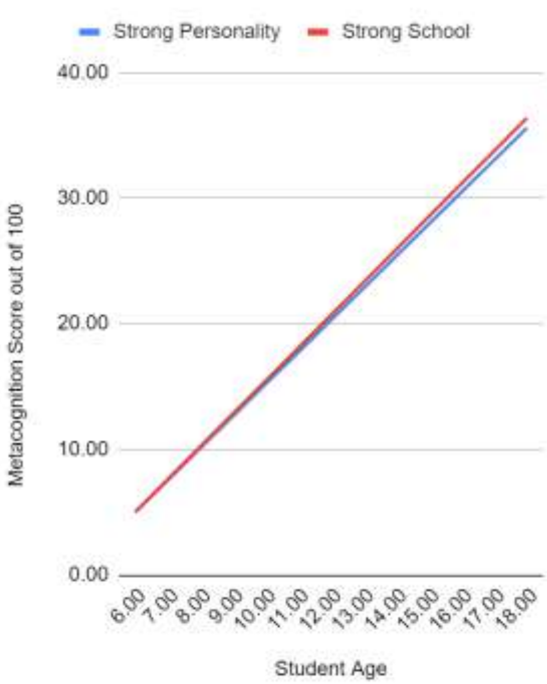

Neuroticism

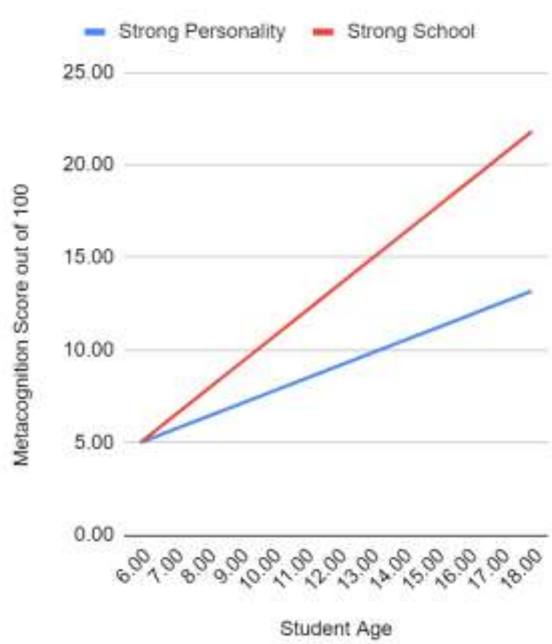

Agreeableness 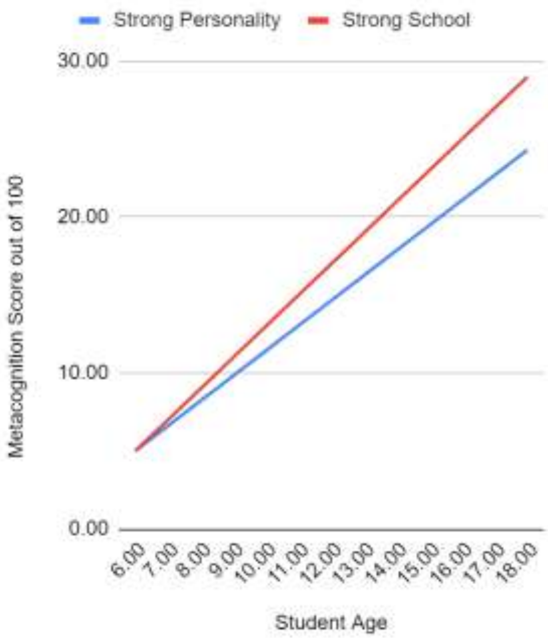	

Table 4: Effectivity of Interventions by Degrees of Correlation

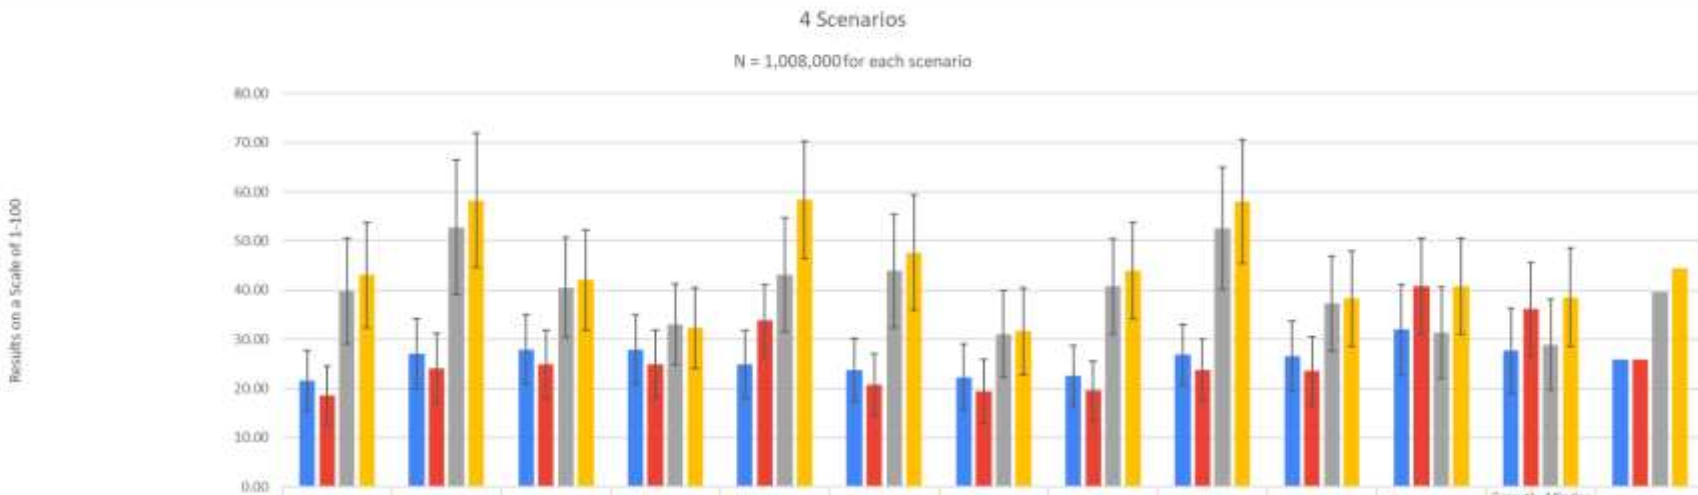

	Creativity	Critical_Thinking	Communication	Collaboration	Mindfulness	Curiosity	Courage	Resilience	Ethics	Leadership	Metacognition	Growth_Mindset	Overall_Average
Mild Correlation, General Intervention	21.59	27.07	27.93	27.93	24.90	23.78	22.27	22.53	26.81	26.60	31.95	27.63	25.91
Mild Correlation, Meta Intervention	18.64	24.07	24.94	24.94	33.89	20.79	19.39	19.53	23.81	23.61	40.73	36.10	25.87
Strong Correlation, General Intervention	39.82	52.80	40.51	32.98	43.10	43.87	31.10	40.74	52.54	37.30	31.35	28.94	39.59
Strong Correlation, Meta Intervention	43.08	58.28	42.05	32.27	58.30	47.67	31.61	43.96	58.02	38.27	40.73	38.56	44.40

Appendix C: Learning Analytics Process Model

Our approach is informed by the Learning Analytics Process Model from Jim Goodell and Steve Ritter. Specifically, we used the Predict, Infer, and Mine approach to come up with what the large dataset might look like, observe the results to see what the data could be telling us, and then use machine learning to confirm those predictions.
Here is our process:

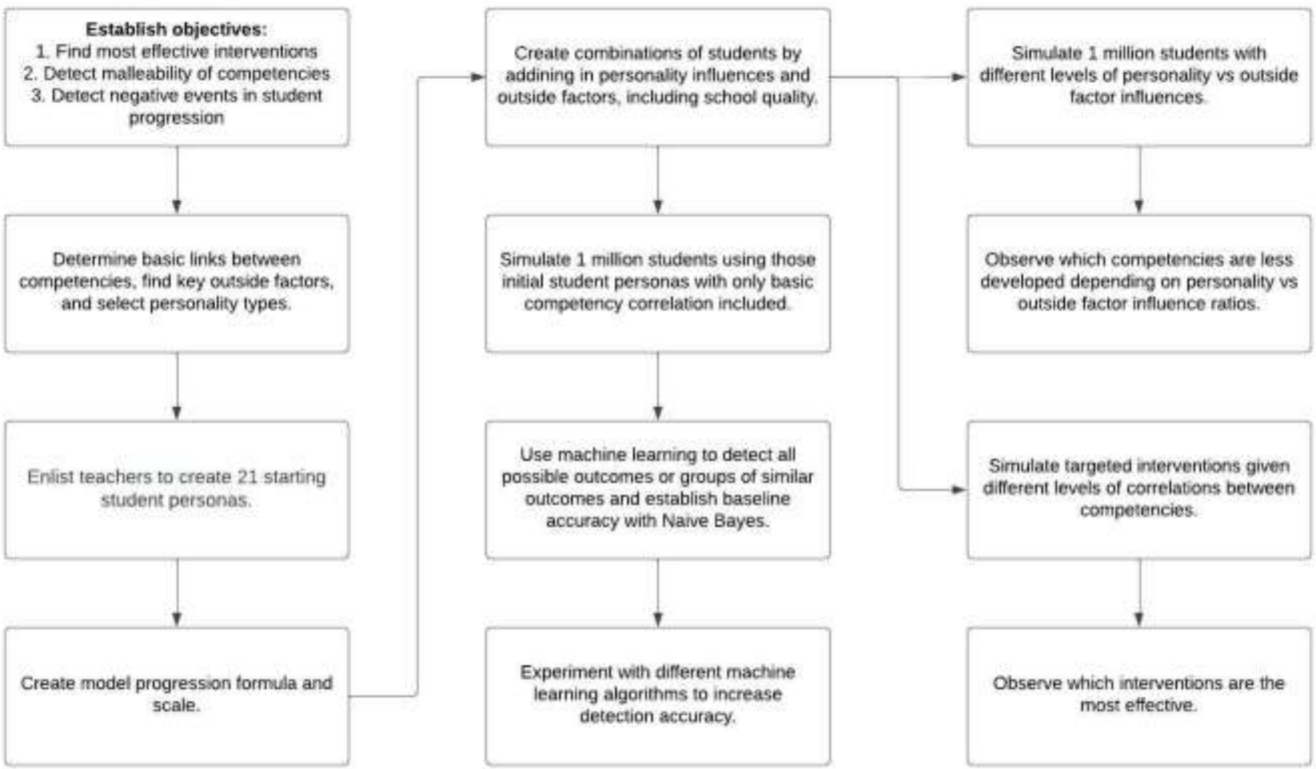

The Learning Analytics Process Model from Goodell and Ritter:[1]

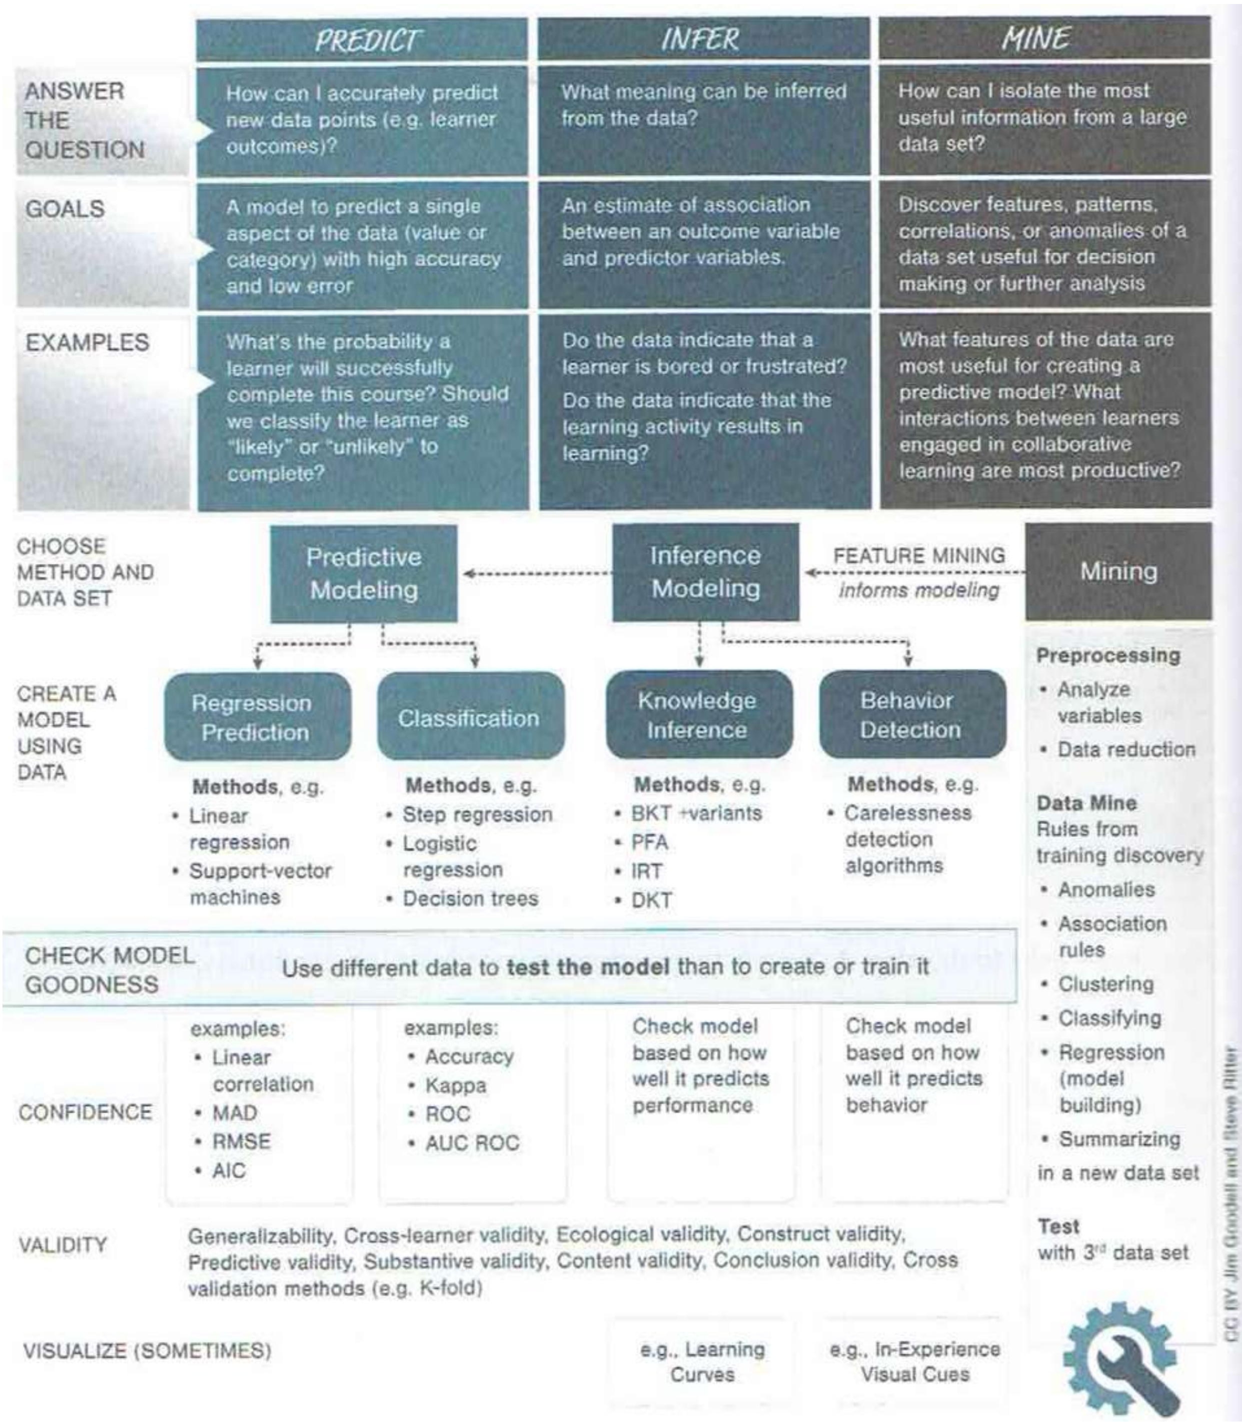

[1] ("Learning Engineering Toolkit" (Routledge, 2022) by Jim Goodell et al. p.366, fig. 18.1 by Goodell and Ritter.)