


**EXAME DE PROFICIÊNCIA EM LÍNGUA INGLESA
MESTRADO / DOUTORADO EM LINGUÍSTICA 2023 – 2º SEMESTRE**

CANDIDATO(A) No: _____

ATENÇÃO!

1. As respostas deverão ser dadas em **português** e **não devem ser uma tradução literal do texto**.
2. As respostas deverão ater-se exclusivamente ao texto e ao que está sendo perguntado.
3. O limite de espaço para cada resposta é de meia página.
4. As respostas deverão ser escritas à tinta.
5. É facultado o uso de dicionários bilíngues e/ou monolíngues, que não poderão ser compartilhados entre os candidatos.
6. O tempo de duração da prova é de 3 (três) horas.

Leia o texto “**How Duolingo’s AI learns what you need to learn**” de autoria de Klinton Bocknell, Claire Brust e Burr Settles, extraído de IEEE SPECTRUM, March 2023, e responda as questões a seguir:



HOW
DUOLINGO'S
AI LEARNS
WHAT
YOU NEED
TO LEARN

The language-learning app tries to emulate a great human tutor

BY KLINTON BICKNELL,
CLAIRE BRUST
& BURR SETTLES



IT'S LUNCHTIME WHEN

your phone pings you with a green owl who cheerily reminds you to “Keep Duo Happy!” It’s a nudge from Duolingo, the popular language-learning app, whose algorithms know you’re most likely to do your 5 minutes of Spanish practice at this time of day. The app chooses its notification words based on what has worked for you in the past and the specifics of your recent achievements, adding a dash of attention-catching novelty. When you open the app, the lesson that’s queued up is calibrated for your skill level, and it includes a review of some words and concepts you flubbed during your last session.

Duolingo, with its gamelike approach and cast of bright cartoon characters, presents a simple user interface to guide learners through a curriculum that leads to language proficiency, or even fluency. But behind the scenes, sophisticated artificial-intelligence (AI) systems are at work. One system in particular, called Birdbrain, is continuously improving the learner’s experience with algorithms based on decades of research in educational psychology, combined with recent advances in machine learning. But from the learner’s perspective, it simply feels as though the green owl is getting better and better at personalizing lessons.

The three of us have been intimately involved in creating and improving Birdbrain, of which Duolingo recently launched its second version. We see our work at Duolingo as furthering the company’s overall mission to “develop the best education in the world and make it universally available.” The AI systems we continue to refine are necessary to scale the learning experience beyond the more than 50 million active learners who currently complete about 1 billion exercises per day on the platform.

Although Duolingo is known as a language-learning app, the company’s ambitions go further. We recently launched apps covering childhood literacy and third-grade mathematics, and these expansions are just the beginning. We hope that anyone who wants help with academic learning will one day be able to turn to the friendly green owl in their pocket who hoots at them, “Ready for your daily lesson?”



BACK IN 1984, educational psychologist Benjamin Bloom identified what has come to be called Bloom’s 2-sigma problem. Bloom found that average students who were individually tutored performed two standard deviations better than they would have in a classroom. That’s enough to raise a person’s test scores from the 50th percentile to the 98th.

When Duolingo was launched in 2012 by Luis von Ahn and Severin Hacker out of a Carnegie Mellon University research project, the goal was to make an easy-to-use online language tutor that could approximate that supercharging effect. The founders weren’t trying to replace great teachers. But as immigrants themselves (from Guatemala and Switzerland, respectively), they recognized that not everyone has access to great teachers. Over the ensuing years, the growing Duolingo team continued to think about how to automate three key attributes of good tutors: They know the material well, they

keep students engaged, and they track what each student currently knows, so they can present material that’s neither too easy nor too hard.

Duolingo uses machine learning and other cutting-edge technologies to mimic these three qualities of a good tutor. First, to ensure expertise, we employ natural-language-processing tools to assist our content developers in auditing and improving our 100-odd courses in more than 40 different languages. These tools analyze the vocabulary and grammar content of lessons and help create a range of possible translations (so the app will accept learners’ responses when there are multiple correct ways to say something). Second, to keep learners engaged, we’ve gamified the experience with points and levels, used text-to-speech tech to create custom voices for each of the characters that populate the Duolingo world, and fine-tuned our notification systems. As for getting inside learners’ heads and giving them just the right lesson—that’s where Birdbrain comes in.

Birdbrain is crucial because learner engagement and lesson difficulty are related. When students are given material that’s too difficult, they often get frustrated and quit. Material that feels easy might keep them engaged, but it doesn’t challenge them as much. Duolingo uses AI to keep its learners squarely in the zone where they remain engaged but are still learning at the edge of their abilities.

One of us (Settles) joined the company just six months after it was founded, helped establish various research functions, and then led Duolingo’s AI and machine-learning efforts until last year. Early on, there weren’t many organizations doing large-scale online interactive learning. The closest analogue to what Duolingo was trying to do were programs that took a “mastery learning” approach, notably for math tutoring. Those programs offered up problems around a similar concept (often called a “knowledge component”) until the learner demonstrated sufficient mastery before moving on to the next unit, section, or concept.

But that approach wasn’t necessarily the best fit for language, where a single exercise can involve many different concepts that interact in complex ways (such as vocabulary, tenses, and grammatical gender), and where there are different ways in which a learner can respond (such as translating a sentence, transcribing an audio snippet, and filling in missing words).



The early machine-learning work at Duolingo tackled fairly simple problems, like how often to return to a particular vocabulary word or concept (which drew on educational research on spaced repetition). We also analyzed learners' errors to identify pain points in the curriculum and then reorganized the order in which we presented the material.

Duolingo then doubled down on building personalized systems. Around 2017, the company started to make a more focused investment in machine learning, and that's when coauthors Brust and Bicknell joined the team. In 2020, we launched the first version of Birdbrain.



BEFORE BIRDBRAIN, Duolingo had made some non-AI attempts to keep learners engaged at the right level, including estimating the difficulty of exercises based on heuristics such as the number of words or characters in a

sentence. But the company often found that it was dealing with trade-offs between how much people were actually learning and how engaged they were. The goal with Birdbrain was to strike the right balance.

Illustration by Eddie Guy

The question we started with was this: For any learner and any given exercise, can we predict how likely the learner is to get that exercise correct? Making that prediction requires Birdbrain to estimate both the difficulty of the exercise and the current proficiency of the learner. Every time a learner completes an exercise, the system updates both estimates. And Duolingo uses the resulting predictions in its session-generator algorithm to dynamically select new exercises for the next lesson.

When we were building the first version of Birdbrain, we knew it needed to be simple and scalable, because we'd be applying it to hundreds of millions of exercises. It needed to be fast and require little computation. We decided to use a flavor of logistic regression inspired by item-response theory from the psychometrics literature. This approach models the probability of a person giving a correct response as a function of two variables, which can be interpreted as the difficulty of the exercise and the ability of the learner. We estimate the difficulty of each exercise by summing up the difficulty of its component features like the type of

exercise, its vocabulary words, and so on.

The second ingredient in the original version of Birdbrain was the ability to perform computationally simple updates on these difficulty and ability parameters. We implement this by performing one step of stochastic gradient descent on the relevant parameters every time a learner completes an exercise. This turns out to be a generalization of the Elo rating system, which is used to rank players in chess and other games. In chess, when a player wins a game, their ability estimate goes up and their opponent's goes down. In Duolingo, when a learner gets an exercise wrong, this system lowers the estimate of their ability and raises the estimate of the exercise's difficulty. Just like in chess, the size of these changes depends on the pairing: If a novice chess player wins against an expert player, the expert's Elo score will be substantially lowered, and their opponent's score will be substantially raised. Similarly, here, if a beginner learner gets a hard exercise correct, the ability and difficulty parameters can shift dramatically, but if the model already expects the learner to be correct, neither parameter changes much.

1. Com base na leitura do título e dos dois primeiros parágrafos do artigo, identifique o assunto do texto e o sistema descrito. (2)

Referência de resposta: O texto aborda um sistema de aprendizagem de língua estrangeira por meio de um *app* denominado *Duolingo*, que incorpora um sofisticado mecanismo de inteligência artificial, denominado *Birdbrain*, o qual se baseia em algoritmos, para aprimorar a experiência de aprendizagem do usuário. Dessa forma, cada vez que ele abre o *app*, as aulas são adequadas ao seu nível de conhecimento, além de incorporar uma revisão de palavras e conceitos anteriormente estudados.

2. Qual é o principal objetivo do sistema e suas perspectivas de desenvolvimento no futuro? (2)

Referência de resposta: O sistema tem como objetivo oferecer educação a um número cada vez maior de usuários, que estão hoje em torno de mais de 50 milhões de alunos ativos. Esses usuários desenvolvem cerca de um bilhão de exercícios diariamente na plataforma, que recentemente incorporou também a alfabetização e a matemática como recursos. No futuro, os desenvolvedores esperam que todos que necessitem de ajuda com seu aprendizado acadêmico possam lançar mão do *app*.

3. Quais são os 3 (três) principais atributos do sistema e como ele os opera? (3)

Referência de resposta: Em primeiro lugar, os desenvolvedores conhecem bem o material e usam a aprendizagem por máquina e outras tecnologias de ponta na elaboração de ferramentas de processamento de linguagem natural para assistir os profissionais na confecção de materiais para cerca de 100 cursos em mais de 40 idiomas. Essas ferramentas analisam o conteúdo vocabular e gramatical das unidades e selecionam um conjunto de traduções possíveis para nortear as respostas dos usuários. Em segundo lugar, os criadores mantêm os alunos envolvidos nas tarefas por meio da gamificação das atividades, com pontuação e níveis, tecnologia que transfere texto e voz e um sistema de notificações. Em terceiro lugar, o sistema mapeia o conhecimento adquirido pelos usuários por meio de inteligência artificial, “calibrando” os conteúdos que serão apresentados a cada novo acesso, de forma a equilibrar o limite entre a manutenção do interesse e novos aprendizados.

4. Como o autor do texto compara e contrasta o sistema para aprendizagem de línguas com outros sistemas, como para aprendizagem de matemática? Discuta (3)

Referência de resposta: O sistema de tutoria em matemática desenvolveu-se a partir de um conceito análogo ao do Duolingo, que seria oferecer problemas em torno de um dado conceito, o que denominaram um “componente do conhecimento”, até que os usuários o dominem e possam passar para a próxima unidade, seção, ou conceito. No entanto, os autores observam que tal abordagem não necessariamente se aplica ao aprendizado de língua, visto que um dado exercício pode envolver vários conceitos que interagem entre si de formas complexas (como vocabulário, tempo verbal e gênero gramatical) e que podem gerar respostas de ordem distinta também, como traduzir uma sentença, transcrever um áudio ou preencher lacunas.
