非官方中文译文：The Bitter Lesson，Rich Sutton（理查德·萨顿），2019-03-13 发表。原文：
http://www.incompleteideas.net/IncIdeas/BitterLesson.html
翻译方式：用浏览器插件「沉浸式翻译」在原文页面上逐段机器翻译后保存，译者只检查了排版，没有逐句校对。 原文采集时间：2026-10-27 13:41:31+00:00。译文版本：v1.0。
译者与原作者没有关联，译文未经作者审阅，以原文为准；译者不对译文的准确性负责，读者据此作出的判断与行动后果自负。原文著作权归原作者所有，下架请求：leoycwan@gmail.com

---

The Bitter Lesson 苦涩的教训

Rich Sutton 理查德·萨顿

March 13, 2019 2019 年 3 月 13 日

The biggest lesson that can be read from 70 years of AI research is that general methods that leverage computation are ultimately the most effective, and by a large margin. The ultimate reason for this is Moore's law, or rather its generalization of continued exponentially falling cost per unit of computation. Most AI research has been conducted as if the computation available to the agent were constant (in which case leveraging human knowledge would be one of the only ways to improve performance) but, over a slightly longer time than a typical research project, massively more computation inevitably becomes available. Seeking an improvement that makes a difference in the shorter term, researchers seek to leverage their human knowledge of the domain, but the only thing that matters in the long run is the leveraging of computation. These two need not run counter to each other, but in practice they tend to. Time spent on one is time not spent on the other. There are psychological commitments to investment in one approach or the other. And the human-knowledge approach tends to complicate methods in ways that make them less suited to taking advantage of general methods leveraging computation. There were many examples of AI researchers' belated learning of this bitter lesson, and it is instructive to review some of the most prominent.

从 70 年的人工智能研究中可以得出的最大教训是：利用计算能力的通用方法最终最为有效，而且优势远远超过其他方法。这背后的根本原因是摩尔定律，或者更准确地说，是其推广形式：单位计算成本持续呈指数级下降。大多数人工智能研究都仿佛假设智能体可用的计算能力是固定不变的（在这种情况下，利用人类知识几乎是提升性能的唯一途径），但只要时间稍长于一个典型研究项目，可用的计算能力就必然会大幅增长。为了在较短时间内取得显著进步，研究人员力求利用自己对领域的知识，但从长远来看，唯一重要的是对计算能力的利用。这两种做法本不必相互冲突，但实践中往往如此。花在其中一种做法上的时间，就无法花在另一种上。人们在心理上会倾向于投入其中一种方法，而非另一种。利用人类知识的方法还往往会使方法变得复杂，导致它们不太适合利用计算能力的通用方法。人工智能研究人员迟迟才领悟这一苦涩教训的例子不胜枚举，回顾其中一些最为突出的案例，颇有启发。

In computer chess, the methods that defeated the world champion, Kasparov, in 1997, were based on massive, deep search. At the time, this was looked upon with dismay by the majority of computer-chess researchers who had pursued methods that leveraged human understanding of the special structure of chess. When a simpler, search-based approach with special hardware and software proved vastly more effective, these human-knowledge-based chess researchers were not good losers. They said that ``brute force'' search may have won this time, but it was not a general strategy, and anyway it was not how people played chess. These researchers wanted methods based on human input to win and were disappointed when

they did not.

在计算机国际象棋领域，1997 年击败世界冠军卡斯帕罗夫的方法，建立在大规模、深层搜索之上。当时，大多数计算机国际象棋研究人员对此深感失望，因为他们一直致力于利用人类对国际象棋特殊结构的理解来设计方法。当一种更简单、以搜索为基础并配合专用软硬件的方法被证明效果好得多时，这些依赖人类知识的国际象棋研究人员并没有坦然接受失败。他们说，“蛮力”搜索这次或许赢了，但它并非一种通用策略，而且无论如何，这并不是人们下棋的方式。这些研究人员希望基于人类输入的方法获胜；当它们没能获胜时，他们感到失望。

A similar pattern of research progress was seen in computer Go, only delayed by a further 20 years. Enormous initial efforts went into avoiding search by taking advantage of human knowledge, or of the special features of the game, but all those efforts proved irrelevant, or worse, once search was applied effectively at scale. Also important was the use of learning by self play to learn a value function (as it was in many other games and even in chess, although learning did not play a big role in the 1997 program that first beat a world champion). Learning by self play, and learning in general, is like search in that it enables massive computation to be brought to bear. Search and learning are the two most important classes of techniques for utilizing massive amounts of computation in AI research. In computer Go, as in computer chess, researchers' initial effort was directed towards utilizing human understanding (so that less search was needed) and only much later was much greater success had by embracing search and learning.

计算机围棋的发展也呈现出类似的研究进步模式，只是又晚了 20 年。最初，人们投入了大量精力，试图利用人类知识或围棋的特殊特性来避免搜索；但一旦搜索得以有效地大规模应用，所有这些努力便都证明无关紧要，甚至适得其反。同样重要的是通过自我对弈进行学习，以学得一个价值函数（其他许多棋类游戏中也是如此，甚至国际象棋也如此，尽管在 1997 年首个击败世界冠军的程序中，学习并未发挥重要作用）。自我对弈学习，以及一般意义上的学习，都和搜索一样，能够让大规模计算发挥作用。在人工智能研究中，搜索和学习是利用海量计算最重要的两类技术。在计算机围棋中，正如在计算机国际象棋中一样，研究人员最初致力于利用人类的理解（以减少所需的搜索量），而直到很久以后，拥抱搜索与学习才带来了大得多的成功。

In speech recognition, there was an early competition, sponsored by DARPA, in the 1970s. Entrants included a host of special methods that took advantage of human knowledge---knowledge of words, of phonemes, of the human vocal tract, etc. On the other side were newer methods that were more statistical in nature and did much more computation, based on hidden Markov models (HMMs). Again, the statistical methods won out over the human-knowledge-based methods. This led to a major change in all of natural language processing, gradually over decades, where statistics and computation came to dominate the field. The recent rise of deep learning in speech recognition is the most recent step in this consistent direction. Deep learning methods rely even less on human knowledge, and use even more computation, together with learning on huge training sets, to produce dramatically better speech recognition systems. As in the games, researchers always tried to make systems that worked the way the researchers thought their own minds worked---they tried to put that knowledge in their systems---but it proved ultimately counterproductive, and a colossal waste of researcher's time, when, through Moore's law, massive computation became available and a means was found to put it to good use.

在语音识别领域，20 世纪 70 年代曾有一场由美国国防高级研究计划局（DARPA）赞助的早期竞赛。参赛者中有许多采用特殊方法的系统，它们利用了人类知识——例如词语、音素、人类声道等方面的知识。另一方则是更新的方法，其本质更偏向统计，所需计算量也大得多，主要基于隐马尔可夫模型（HMM）。再一次，统计方法胜过了基于人类知识的方法。这使整个自然语言处理领域发生了重大转变：在此后的几十年间，统计与计算逐渐占据主导地位。近来，深度学习在语音识别领域的兴起，是这一持续发展方向上的最新一步。深度学习方法对人类知识的依赖更少，同时运用更多计算资源，并在庞大的训练集上进行学习，从而打造出性能大幅提升的语音识别系统。和棋类游戏一样，研究人员总是试图构建出按他们所设想的自身思维方式运作的系统——他们试图把这些知识装进系统里——但事实证明，这最终适得其反；当摩尔定律带来了海量计算能力，且人们找到了善用它的方法后，这种做法便成了对研究人员时间的巨大浪费。

In computer vision, there has been a similar pattern. Early methods conceived of vision as searching for edges, or generalized cylinders, or in terms of SIFT features. But today all this is discarded. Modern deep-learning neural networks use only the notions of convolution and certain kinds of invariances, and perform much better.

在计算机视觉领域，也出现过类似的模式。早期方法将视觉理解为对边缘、广义圆柱体，或 SIFT 特征的搜索。但如今，这些做法都已被摒弃。现代深度学习神经网络只利用卷积以及某些类型的不变性概念，表现却好得多。

This is a big lesson. As a field, we still have not thoroughly learned it, as we are continuing to make the same kind of mistakes. To see this, and to effectively resist it, we have to understand the appeal of these mistakes. We have to learn the bitter lesson that building in how we think we think does not work in the long run. The bitter lesson is based on the historical observations that 1) AI researchers have often tried to build knowledge into their agents, 2) this always helps in the short term, and is personally satisfying to the researcher, but 3) in the long run it plateaus and even inhibits further progress, and 4) breakthrough progress eventually arrives by an opposing approach based on scaling computation by search and learning. The eventual success is tinged with bitterness, and often incompletely digested, because it is success over a favored, human-centric approach.

这是一个重要的教训。作为一个领域，我们至今仍未彻底吸取这一教训，因为我们仍在犯同一类错误。要看清这一点并有效抵制这种倾向，我们就必须理解这些错误为何如此诱人。我们必须吸取这一苦涩的教训：把我们自以为的思维方式嵌入系统，从长远来看行不通。这一苦涩教训源于历史观察：1）人工智能研究者常常试图将知识植入他们的智能体；2）这总能在短期内带来帮助，也让研究者本人感到满意；但 3）从长远来看，这种做法会陷入瓶颈，甚至阻碍进一步进展；4）最终，突破性进展总是通过相反的路径到来，即依靠搜索和学习来扩展计算能力。这种最终的成功带着一丝苦涩，而且往往未被充分消化，因为它战胜的是一种备受青睐、以人为中心的方法。

One thing that should be learned from the bitter lesson is the great power of general purpose methods, of methods that continue to scale with increased computation even as the available computation becomes very great. The two methods that seem to scale arbitrarily in this way are *search* and *learning*.

我们应从这一苦涩教训中汲取的一点是：通用方法具有巨大的力量。即使可用的计算资源变得极其庞大，这些方法仍能随着计算能力的提升而持续扩展。两种似乎能够以这种方式无限扩展的方法是搜索和学习。

The second general point to be learned from the bitter lesson is that the actual contents of minds are tremendously, irredeemably complex; we should stop trying to find simple ways to think about the contents

of minds, such as simple ways to think about space, objects, multiple agents, or symmetries. All these are part of the arbitrary, intrinsically-complex, outside world. They are not what should be built in, as their complexity is endless; instead we should build in only the meta-methods that can find and capture this arbitrary complexity. Essential to these methods is that they can find good approximations, but the search for them should be by our methods, not by us. We want AI agents that can discover like we can, not which contain what we have discovered. Building in our discoveries only makes it harder to see how the discovering process can be done.

从苦涩教训中得到的第二个普遍启示是：心智的实际内容极其复杂，复杂到无可挽回；我们应该停止试图寻找简单的方式来思考心智的内容，比如用简单的方式思考空间、物体、多个智能体或对称性。这些都属于任意且本质上复杂的外部世界。它们不应该被内置，因为它们的复杂性无穷无尽；相反，我们只应内置那些能够发现并捕捉这种任意复杂性的元方法。这些方法的关键在于能够找到良好的近似，但寻找近似的工作应由我们的方法来完成，而不是由我们亲自完成。我们想要的是能够像我们一样进行发现的 AI 智能体，而不是装载着我们已发现成果的智能体。把我们的发现内置进去，只会让人更难看清发现过程该如何实现。

非官方中文译文：The Bitter Lesson，Rich Sutton（理查德·萨顿），2019-03-13 发表。原文：
http://www.incompleteideas.net/IncIdeas/BitterLesson.html
翻译方式：用浏览器插件「沉浸式翻译」在原文页面上逐段机器翻译后保存，译者只检查了排版，没有逐句校对。 原文采集时间：2026-10-27 13:41:31+00:00。译文版本：v1.0。
译者与原作者没有关联，译文未经作者审阅，以原文为准；译者不对译文的准确性负责，读者据此作出的判断与行动后果自负。原文著作权归原作者所有，下架请求：leoycwan@gmail.com