


Feature Extraction

Foundations and Applications

Everyone loves a good competition. As I write this, two billion fans are eagerly anticipating the 2006 World Cup. Meanwhile, a fan base that is somewhat smaller (but presumably includes you, dear reader) is equally eager to read all about the results of the NIPS 2003 Feature Selection Challenge, contained herein. Fans of Radford Neal and Jianguo Zhang (or of Bayesian neural networks and Dirichlet distribution trees) are gloating "I told you so" and looking for proof that their win was not a fluke. But the matter is by no means settled, and fans of SVMs are shouting "wait 'til next year!" You know this book is a bit more edgy than your standard academic treatise as soon as you see the dedication: "To our friends and foes." Competition breeds improvement. Fifty years ago, the champion in 100m butterfly swimming was 22 percent slower than today's champion; the women's marathon champion from just 30 years ago was 26 percent slower. Who knows how much better our machine learning algorithms would be today if Turing in 1950 had proposed an effective competition rather than his elusive Test? But what makes an effective competition? The field of Speech Recognition has had NIST-run competitions since 1988; error rates have been reduced by a factor of three or more, but the field has not yet had the impact expected of it. Information Retrieval has had its TREC competition since 1992; progress has been steady and refugees from the competition have played important roles in the hundred-billion-dollar search industry. Robotics has had the DARPA Grand Challenge for only two years, but in that time we have seen the results go from complete failure to resounding success (although it may have helped that the second year's course was somewhat easier than the first's).

Everyone loves a good competition. As I write this, two billion fans are eagerly anticipating the 2006 World Cup. Meanwhile, a fan base that is somewhat smaller (but presumably includes you, dear reader) is equally eager to read all about the results of the NIPS 2003 Feature Selection Challenge, contained herein. Fans of Radford Neal and Jianguo Zhang (or of Bayesian neural networks and Dirichlet distribution trees) are gloating "I told you so" and looking for proof that their win was not a fluke. But the matter is by no means settled, and fans of SVMs are shouting "wait 'til next year!" You know this book is a bit more edgy than your standard academic treatise as soon as you see the dedication: "To our friends and foes." Competition breeds improvement. Fifty years ago, the champion in 100m butterfly swimming was 22 percent slower than today's champion; the women's marathon champion from just 30 years ago was 26 percent slower. Who knows how much better our machine learning algorithms would be today if Turing in 1950 had proposed an effective competition rather than his elusive Test? But what makes an effective competition? The field of Speech Recognition has had NIST-run competitions since 1988; error rates have been reduced by a factor of three or more, but the field has not yet had the impact expected of it. Information Retrieval has had its TREC competition since 1992; progress has been steady and refugees from the competition have played important roles in the hundred-billion-dollar search industry. Robotics has had the DARPA Grand Challenge for only two years, but in that time we have seen the results go from complete failure to resounding success (although it may have helped that the second year's course was somewhat easier than the first's).



320,99 €

299,99 € (zzgl. MwSt.)

Lieferfrist: bis zu 10 Tage

Artikelnummer: 9783540354871

Medium: Buch

ISBN: 978-3-540-35487-1

Verlag: Springer

Erscheinungstermin: 20.07.2006

Sprache(n): Englisch

Auflage: 1. Auflage 2006

Serie: Studies in Fuzziness and Soft Computing

Produktform: Gebunden

Gewicht: 1364 g

Seiten: 778

Format (B x H): 155 x 235 mm

